0

RAND

An Approach to Understanding the Value of Parts

Marygail K. Brauner, James S. Hodges, Daniel A. Relles

Arroyo Center

The research described in this report was sponsored jointly by the Navy Secretariat, Naval Air Systems Command (NAVAIR-43), Naval Supply System (NAVSUP), and the Navy's Aviation Suppy Office (ASO).

Library of Congress Cataloging in Publication Data

Brauner, Marygail K., 1947–
An approach to understanding the value of parts / Marygail K. Brauner, James S. Hodges, and Daniel A. Relles.
p. cm.
"Prepared for the United States Navy."
"MR-313-A/USN."
Includes bibliographical references.
ISBN 0-8330-1509-5
1. United States. Navy—Aviation supplies and stores.
2. United States. Navy—Equipment—Maintenance and repair.
3. United States. Navy—Cost control. I. Hodges, James S.
II. Relles, Daniel A. III. United States. Navy. IV. RAND.
V. Title.
VG93.B717 1994
359.9′48′0973—dc20 94-3007
CIP

Published 1994 by RAND
1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
To obtain information about RAND studies or to order documents,
call Distribution Services, (310) 451-7002

RAND

An Approach to Understanding the Value of Parts

Marygail K. Brauner, James S. Hodges, Daniel A. Relles

Prepared for the United States Navy

Arroyo Center

Approved for public release; distribution unlimited

PREFACE

This work is part of a larger RAND project entitled "Enhancing the Logistics System: The Depot Perspective," sponsored jointly by the Navy Secretariat, Naval Air Systems Command (AIR-43), Naval Supply System, and Aviation Supply Office. This project, which began in 1989, has three objectives:

- Improve the readiness and sustainability of Naval aviation
- Improve the integration of Naval aviation logistics
- Identify cost-reduction opportunities.

The RAND report *Improving Naval Aviation Depot Responsiveness* (R-4133-A/USN, 1992, by M. K. Brauner, D. A. Relles, and L. A. Galway) documents some of the work completed on the first two objectives. This report presents work on the third objective.

Our mandate in the project includes suggesting and analyzing broad policy issues, as well as offering more detailed suggestions. RAND work on the broader policy issues is described in the aforementioned report and in two additional RAND documents: *Materiel Problems at a Naval Aviation Depot: A Case Study of the TF-30 Engine* (N-3473-A/USN, 1992) and *Management Adaptations in Jet Engine Repair at a Naval Aviation Depot in Support of Operation Desert Shield/Storm* (N-3436-A/USN, 1992), both by L. A. Galway.

This report develops a model of the repair pipeline that leads to an algorithm for stocking parts according to the parts' ability to reduce the value of the repair pipeline. It then uses available Navy data to develop and evaluate various stockage policies through a simulation.

Naval maintenance and supply officers will find this work of interest. Since some of the problems faced by the Naval aviation logistics system are common to all services and the private sector, this work will also be of interest to logisticians in the other services, in DoD, and in commercial remanufacturing.

CONTENTS

FIGURES

TABLES

SUMMARY

PURPOSE

The objective of the Naval aviation logistics system is to keep Naval aircraft flying. That system has many parts, of which the Naval aviation depots (NADEPs) are but one. NADEPs maintain aircraft and repair some of their broken components. A key constraint on NADEP activities is materiel support—ensuring the availability of parts. Absence of key repair parts blocks the execution of repair schedules designed to meet the needs of the fleet, and, hence, has a direct effect on readiness.

This report describes an approach to materiel management that has as a goal reduced NADEP turnaround time, as well as less excess materiel in the system.

THE REPAIR-PARTS PROBLEM

Forecasting—regardless of model—is never perfect. And, for aviation repairs at least, there is too much uncertainty in demands. The actual performance of a part depends on operating conditions, age of the part, quality of manufacturing, frequency of use, etc. Shortages and excesses are inevitable, but dealing with them is not simply a matter of forecasting demand better, then procuring more efficiently to the forecast.

For planners to be able to improve repair facility responsiveness, the repair-parts problem must be defined in terms of factors that planners can actually improve. Such a definition is *ensuring that parts are where they are needed, when they are needed.* With this definition, the repair-parts problem can be reduced to two elements: the parts requirement, and the time to get the parts. Repair facilities know how to specify parameters of their parts requirements. They maintain bills of materials (BOM) for end-items and replacement factors (RFs) for the individual parts on the end-item, and they regularly update these bits of information. The time-to-get portion of the problem, however, is more complicated. It is a combination of setting authorized stockage levels, knowing where parts are located, and knowing how long it will take for the repair facility to receive off-facility parts.

We attempt here to understand the time-to-get portion of the problem. Our research leads to the definition of the *value of a part,* where *value* is essentially how much the part contributes to shortening the repair time of the end-item that uses it. Regardless of what a part costs, if the part breaks frequently and tends to hold up the repair of an expensive end-item, then the part is very valuable. Conversely, if a part is never needed, its value is zero.

Knowing this value of a repair part can help planners improve repair facility responsiveness. Because many planners involved in remanufacturing (e.g., those repairing aircraft engines or complex electronic components) must rely on expensive repairable parts, a measure of value can help them make appropriate decisions about what parts to stock or to invest in. Our suggestion is to set authorized stockage levels for repairable and consumable parts to minimize the expected dollar value of the repair pipeline for a given level of investment. In this document, this method is described narratively; the mathematics are provided in a companion document, *Models and Algorithms for Repair Parts Investment and Management*, MR-314-A/USN, 1993, by James S. Hodges.

STOCKING PARTS TO MAXIMIZE VALUE FOR A GIVEN INVESTMENT

We seek to quantify the cost—in terms of both time and money—to the logistics system of not having readily available repair parts. To do so, we develop a stockage method

that incorporates stochastic (random) demands and repairs[1] and that describes the flow of repairable end-items (called weapon-replaceable assemblies [WRAs] in the Navy) through the depots. Our measure of value is essentially an estimate of the return on investment for a given part.

In our computations, we used data the Navy routinely collects—Navy Inventory Materiel Management System (NIMMS) data from the NADEPs and price information from the Aviation Supply Office (ASO). The former data describe which end-items were repaired during a three-year period, what parts were used to repair them, and how long the end-items waited for the parts. The latter data provide each end-item's unit price and help us link up information from disparate sources. We had no source for current inventories and were unable to obtain that kind of information. This lack prevents us from making predictive statements about the effects of alternative policies on the size of the current repair pipeline, although we do attempt to impute current stock and to compare additional stock against imputed stock.

Using our measure of value, we develop a rank-ordered list of repair parts: The higher it is on the list, the more valuable the part is to reducing the value of the repair pipeline. The algorithm to stock parts for a given investment is uncomplicated:

- Step 1. Begin with zero authorized stock for all parts.
- Step 2. For each part, compute the value of the next unit of the part (we discuss how to do this in Section 3) and divide its value by its unit price to obtain the rate of return from stocking that unit.
- Step 3. Select the part with the highest rate of return, and stock one unit of it.
- Step 4. Is the total cost of authorized stock less than the total investment to be made? If so, return to Step 2; if not, stop.

The stockage problem we have posed is an instance of the classic knapsack problem. Our algorithm is a heuristic solution, a so-called greedy algorithm, which has been extensively studied. This greedy algorithm does not, in general, maximize the objective function subject to the constraints, but it is known to produce good solutions as long as the costs of the individual increments to stock are not large relative to the budget constraint. We note that other formulations of the optimization problem are possible, such as an infinite-stage dynamic program with an average reward criterion.[2] For simplicity, we did not use this formulation.

EVALUATING THE METHOD OF BUILDING AUTHORIZED STOCK

We used computer simulation to test whether specific aspects of the mathematical methods introduce error or inefficiency. We demonstrate the promise of the method by first constructing a baseline case that imputes current authorized stockage levels at Navy depots, by then constructing a treatment case that increments the baseline stockage levels (using the method described in Section 4), and by finally computing the difference in performance between the baseline stockage levels and the treatment stockage levels as predicted by our stockage method and as assessed by our simulation tests.

---

[1] The following quotation from *Onward Through the Fog: Uncertainty and Management Adaptation in Systems Analysis and Design* by James S. Hodges (Santa Monica, Calif.: RAND, R-3760-AF/A/OSD, p. 1) explains why a responsive logistics system must account for the erratic behavior involved in component repair:

"If events were predictable, policy analysis would be simple. For example, if Air Force planners knew which aircraft parts would fail and where and when they would fail, it would be straightforward to schedule repair of parts, to buy parts and repair capability, and to make budgets. But nobody can accurately predict which parts will fail, when they will fail, or where they will fail, even in peacetime. The wars for which we prepare are also shrouded in uncertainty."

[2] See, for example, S. M. Ross, *Introduction to Stochastic Dynamic Programming*, New York: Academic Press, 1983, Chapter V.

The tests show that our method does a good job of setting authorized stockage levels. The simulations also suggest that large savings may be possible, and they identify the weapon systems, by special materiel identification code, for which savings are likely to accrue. The results make the case for experimenting with our method at a depot or at a remanufacturing site.

USES OF THE VALUE MEASURE AND EXTENSIONS OF THE METHOD

The method developed in this report produces a rank-ordered list of authorized parts to support a repair process, something EOQ[3] rules tend not to do. Not having the parts to repair an end-item can significantly impede the timely production of remanufactured items. In lean times when budgets are tight, it is important to know which parts to buy first. There are several other ways to use the value of parts idea:

1. Should the Navy decide to stock parts at NADEPs the way it stocks parts on deploying carriers in AVCALs (Aviation Consolidated Allowance Lists), our method could serve as a starting point for building those stocks. The Navy has called this idea a *NADEPCAL*.
2. Value of parts could be used to evaluate current reorder rules for existing stocks and to identify how those rules should change as weapon systems age, the force is drawn down, NADEP repair procedures change, etc.
3. The calculations can be used to balance investment strategies between spending money on parts and spending it on other segments of the repair pipeline: parts reliability, faster requisition processing, faster distribution, improved repair processes, etc.
4. The calculations can be used to identify parts to watch: parts that are problematic now or that may become problematic later.
5. A similar computation of value measure can be used to attribute value to supply actions (e.g., speedup of delivery of due-in items) in terms of the effect such actions have on the availability of aircraft at the end of a specific time horizon.

This report has concentrated on using a value measure to make investments in repair parts to minimize the expected value of the repair pipeline for a given level of investment; in addition, as suggested above in the fifth point, a similar computation of value can be used to help managers make operating decisions that affect the short-run output of the NADEP.[4] With knowledge of the near-term repair schedule, the stock on hand, and the parts due in from supply, managers at NADEPs can take actions to alter the parts due in to service the repair schedule more expeditiously. Using a value measure in this way provides the NADEP with a tool that indicates when it is worthwhile to selectively pay more for better supply service (for example, using Federal Express), targeting expenditures to specific items that will yield the largest improvement in aircraft availability. We have not had the opportunity to test the short-run value measure as we tested the long-run measure.

Experiences in the private sector show that changes in both policies and processes are necessary to achieve order-of-magnitude improvements in remanufacturing operations. In addition to setting authorized stock levels for the NADEPs, the Navy must give attention to improving maintainability and contracting, updating replacement factors and demand rates, and decreasing OST and process times. We conclude that assessing these various options through their effect on the value of the repair pipeline is a good way to proceed.

---

[3]The military services usually set stockage levels based on an economic order quantity (EOQ). The formulas vary, but most include the demand for the item, the cost of holding the inventory, and the cost of processing the order.

[4]The mathematics that underlie the short-run use of the value measure are provided in James S. Hodges, *Models and Algorithms for Repair Parts Investment and Management*, Santa Monica, Calif.: RAND, MR-314-A/USN, 1993.

ACKNOWLEDGMENTS

While developing the approach in this report, we conducted several internal seminars at RAND and received helpful suggestions from Jack Abell, Frank Camm, John Folkeson, Don Gaver, Jean Gebman, Lou Miller, Nancy Moore, Ray Pyles, Tim Ramey, and Hy Shulman. Ken Girardini reviewed the report and made many useful comments. Paul Steinberg and Marian Branch helped make this report less user-hostile.

Over the course of RAND's work in repair parts, we have received generous help from countless civilian and uniformed Navy personnel. Most important to this report has been support from NAVAIR-43, in particular its commander, ADM Don Eaton (ret.), as well as our action officers, CDR Brian Chandler (ret.), CAPT Ryan Hansen, and CAPT Bambi Smith. We have also received much-appreciated help from CAPT Glenn Downer (ret.) at the Navy Secretariat; ADM James Eckelberger (ret.) and ADM James Davidson at the Aviation Supply Office; and CAPT Charles Sapp, LCDR Nick Zimmon, and LCDR Stan Pyle, at NADEP, North Island.

This work would not have been possible without the help of these people. Nonetheless, all assertions and interpretations in this report belong to the author alone.

ACRONYMS AND ABBREVIATIONS

a/c	Aircraft
ALC	Air Logistics Center
ASL	Authorized stockage level
ASO	Aviation Supply Office
AVCAL	Aviation Consolidated Allowance List
AWG-9	Radar system for F-14 aircraft
AWP	Awaiting parts
BOM	Bill of materials
COG	Cognizance code
DISC	Defense Industrial Supply Center
DLA	Defense Logistics Agency
DOCNO	Document number
DoD	Department of Defense
EOQ	Economic order quantity
ICP	Inventory control point
IIC	Item identification code
NADEP	Naval aviation depot
NADEPCAL	*Naval Aviation Depot Consolidated Allowance List*
NFMSO	Navy Fleet Materiel Support Office
NIF	Naval Industrial Fund
NIIN	National item identification number
NIMMS	Navy Inventory Materiel Management System
NRFI	Not ready for issue
NSC	Naval Supply Center
OST	Order-and-ship time
RF	Replacement factor
RFI	Ready for issue
ROI	Return on investment
SD	Standard deviation
SDLM	Standard depot-level maintenance
SMIC	Special materiel identification code
SPCC	Navy Ships Parts Control Center
SRA	Shop-replaceable assembly
SROI	Simulated return on investment
TpROI	Tall-pole return on investment
UPA	Units per application
WRA	Weapon-replaceable assembly

1. INTRODUCTION

BACKGROUND

The objective of the Naval aviation logistics system is to keep Naval aircraft flying. That system has many parts, of which the Naval aviation depots (NADEPs) are but one. The role of NADEPs is viewed differently in various quarters. Traditionally, NADEPs have been seen within the Navy as performing such functions as aircraft rework (called standard depot-level maintenance—SDLM) and component repair—functions the intermediate-level maintenance personnel at Naval Air Stations or on aircraft carriers could not perform. In this view, NADEPs contribute indirectly to readiness by keeping shelves stocked in Naval Supply Centers. Others see the NADEPs as direct contributors to readiness. In this view, through responsive transportation, repair scheduling, and allocation of ready for issue items (RFI), NADEP activities are organized to meet the needs of the fleet.

Materiel support—ensuring the availability of parts for NADEP activities—is a key constraint on those activities. Absence of key repair parts blocks the execution of repair schedules designed to respond to fleet needs and hence has a direct impact on readiness. Less directly, materiel support problems affect readiness by forcing the system to make large outlays for end-items, e.g., radar antennas or engine modules, to fill long repair pipelines and to stock excessive quantities of piece parts that may become obsolete before they are needed.

The work described in this document was performed for the Navy, but it has broader applications. Organizations that use valuable equipment, such as aircraft, made up of expensive repairable components will face problems similar to those the Navy faces in keeping the equipment fully functional at least cost. Such organizations operate networks of facilities for supply and maintenance of equipment and repairable components. They must invest in the elements of the supply and maintenance system, allocating funds among spare repairable components, repair parts, repair capacity, second-destination transportation,[1] and so on. At any given time, the managers of such facilities must also make operating decisions—which components to induct, which repair parts to order—that affect the performance of the maintenance and supply system and the cost of achieving that performance. This document proposes a way to think about these investments and operating decisions, focusing on spare repairable components and repair parts. From this point on, all examples will be for the Naval aviation logistics system; it is left to the reader to apply the concepts to other remanufacturing situations.

This report describes an approach to materiel management that has as a goal reduced NADEP turnaround time, as well as less excess materiel in the system and more RFI end-items.

CURRENT PROBLEMS

Major changes are in the offing for NADEPs: Under pressure to reduce both costs and repair turnaround times, three NADEPs are scheduled to close and their work is to be consolidated at the remaining three NADEPs.[2] Operationally, this means that the Navy will no longer have dual repair capability on each coast for its major fighter aircraft (F-14s and F/A-18s). If attention is not paid to responsive transportation, repair turnaround times could increase because parts must be moved from one coast to the other. Additionally, the Defense Logistics Agency (DLA) will assume a much greater role in the supply of consumable materiel (piece parts). There is concern in the Navy that DLA will be unable to respond to

the Navy's needs and that current methods are inadequate to identify critical parts before they cause major degradations in readiness.[3]

One of the Navy's key measures of its ability to supply the parts needed by the fleet is NADEP turnaround time, which is constrained by three factors:

1. Labor availability and labor hours required to fix an end-item—usually called a weapon-replaceable assembly (WRA)
2. Test stand availability and the time a part spends on the test stand
3. Delay due to unavailability of shop-replaceable assemblies (SRAs) or piece parts.

Current level-repair schedules for the NADEPs are intended to create a steady flow of work through the depots and allow for stocking parts needed to complete the repair schedule.[4] However, NADEP artisans often do not have the materiel they need to complete even scheduled repairs, not to mention the portion of the workload that is not on the level-repair schedule.[5] They compensate with work-arounds of all types, e.g., cannibalization of other items in work further down the schedule.

GOAL OF THIS RESEARCH

This report is about finding solutions to the problem of unavailability of parts (Item 3 above). We quantify the reduction in NADEP turnaround time that can be achieved by having repair parts readily available. Better availability of parts will reduce NADEP turnaround time, thus making more WRAs available.

To set the repair parts problem in context, we developed a method that incorporates stochastic (random) demands and repairs, and that describes the flow of WRAs through the depots. This method leads to a measure that we call the *value of parts*. Knowledge of value helps to complete the inventory-management picture: The *cost* of a part is readily accessible; the *value* is not. Regardless of what a part costs, if the part is never needed to repair a WRA, its value is *low*; if absence of a part is likely to keep an expensive WRA waiting, the part's value is *high*. Value helps to quantify the return on investment for a given part.

The research to date consists of the development of a method to establish the range and depth of the parts to stock for a specified budget and process parameters, and an attempt to quantify their effect on managing repair parts. Empirical results suggest the size of reductions in turnaround time to be achieved through this stockage policy with an $(S - s, S)$ reorder discipline.[6] Our results support the idea of experimenting to determine whether savings predicted by the method hold up in practice.

We have chosen here to focus on the effects of alternative stockage rules (our method) and disciplines for ordering new parts. We also include a discussion of the potential benefits from speeding up the piece part transportation and distribution pipelines.

ORGANIZATION OF THIS REPORT

Section 2 of this report defines the repair parts problem. Section 3 conceptualizes our method for managing parts, proposes a policy for stocking parts, and discusses how to compute the value of a part. It introduces a key formula—the tall-pole formula—used for calculating the expected time a WRA waits for parts. Section 4 describes the Navy data used in this research and how we built a rank-ordered list of increments to authorized stock using parameters from those data. Section 5 tests our stockage method by simulation studies. The simulation programs we wrote for the study were robust enough to accommodate variations

in other process parameters. Simulation results for speeded up transportation and distribution pipelines are presented in Section 6. Finally, Section 7 discusses the implications of the research and some possible applications of the ideas developed here.

2. WHAT IS THE REPAIR-PARTS PROBLEM?

Many in the Navy view the repair-parts problem as a forecasting and procurement problem: Forecast demand better, then procure more efficiently to the forecast. The result of this focus has been large excesses of parts: poor forecasting, long procurement lead times, and lengthy repair pipelines lead to buying many wrong components.[7] Attention has been focused on improving predictions rather than on shortening repair-pipeline times, especially order-and-ship times (OSTs) and NADEP turnaround time.

WHY THE PROBLEM WILL NOT IMPROVE WITH BETTER FORECASTING

In theory, forecasting the parts required for repairing an end-item should be a straightforward calculation.[8] Simply taking the bill of materials (BOM) for the WRA and the replacement factors[9] for the individual parts on the WRA, along with a desired level of responsiveness (i.e., stock all parts with replacement factors of 0.5 or greater), yields the required parts.[10] But, when a not ready for issue (NRFI) WRA is disassembled at the NADEP, many of the parts with replacement factors of 0.5 or greater will not be broken and some parts not expected to break will be broken. Thus, because the actual performance of a part depends on flying conditions, age of the part, quality of manufacturing, frequency of use, etc., and because forecasting—regardless of the model—is never perfect, some parts that are not needed will be available and others that are needed will be unavailable.

Replacement factors should change over time, but it is not uncommon for those used in requirements calculations to be the numbers given by the manufacturer when the weapon system was initially fielded. In addition, BOMs that should, in theory, be stable, change periodically as the aircraft is modified.. Until the WRA is inducted at the NADEP and is inspected, there is no knowledge of what modifications have been completed on a particular unit of the WRA.

In addition to variation in the replacement factors and BOMs, there is also considerable variation in the repair schedule. NADEPs would like to plan to a level schedule, for example, a guaranteed 15 AWG-9 radars every quarter. But, in some quarters, fewer than the expected number of WRAs will require repair and, in others, more—because the repair schedule is driven by the flying program and repair capability. The past flying program produced the current mix of NRFI WRAs and SRAs, and the near-term future flying program will produce more broken parts, some of which will need to be repaired quickly to get a downed aircraft flying again. Many WRAs are directly influenced by the types of sorties flown and their frequency. Avionics components seem to benefit from frequent use and deteriorate with periods of idleness. Engine repair is often driven solely by the hours logged. In recent years, flying hours have been relatively stable. But, as the force size is reduced and money for peacetime operations declines, they can be expected to change. It will be difficult to predict the effects of these differences on the demands for repair at NADEPs.

Forecasting over long horizons is very difficult; yet, a long planning horizon is considered in scheduling NADEP repair. Such long horizons are the result of long OSTs, long procurement lead times, and lengthy repair times in the NADEPs.

Uncertainty in demands cannot be reduced significantly. Therefore, focusing on improved forecasting models will have only small effects toward improving NADEP responsiveness.

IF NOT FORECASTING, WHAT?

It would be helpful to attempt to define the repair-parts problem in term of factors that the Navy can improve. Such an alternative definition is *ensuring that parts are where they are needed when they are needed.* With this definition, the repair-parts problem can be reduced to two elements: the parts requirement and the time to get the parts.

Knowing Parts Requirements

Regular updating of BOMs and replacement factors would help NADEPs perform their jobs more efficiently. Galway noted that without heroic efforts by NADEP personnel, critical time would have been lost during the Desert Shield engine-repair surge because no current BOM existed for the T64 engine.[11] Yet this engine had been repaired by the Navy for many years. While the Navy is making progress in developing BOMs and replacement factors, for this research we developed our own BOMs and replacement factors using Navy Inventory Materiel Management System (NIMMS) data (see Section 4).

Getting Parts Quickly

Improving repair-parts availability requires knowing where parts are located and how long it will take for the NADEP to receive off-station parts. In a study of the TF30 engine, for example, some inventory control points (ICPs) were able to provide parts more quickly than others.[12] It was not clear whether the difference was the result of actions at the ICP or of the type of parts the ICP carried. In general, commonly used items that have commercial equivalents and thus are more easily procured will be supplied to the NADEPs more readily than unique parts. Parts stockage must take lead times into account: Parts that can be obtained quickly will require a smaller stock than those with long lead times.

Improving NADEP and Supply System Interaction

There is a further complication to the "time to get" portion of the repair-parts problem that involves how the NADEP interacts with the rest of the logistics system. Although the structure of logistics support in both the Navy and DoD is changing, a basic disconnect between the ICP and NADEP that inhibits NADEP repair and rework activities will probably not change. An example illustrates the disconnect.

Defense Logistics Agency ICPs fill requisitions for consumables out of their own warehouses. In the Navy, these requisitions come from end-users or from Naval Supply Centers (NSCs), which stock DLA-managed items for NADEPs and other customers. DLA ICPs procure and stock consumables solely on the basis of the history of requisitions to DLA, and different stockage rules are used for frequently and infrequently requisitioned items. This DLA policy can create a deleterious feedback loop. For example, a NADEP might need 2 units of an item per year, and order 2 units. The NSC passes on the order to the DLA ICP, which ships that item only in lots of 10 units, but must buy from the vendor in lots of 100. The ICP ships 10 units to the NADEP.

The NADEP uses those 10 units over five years, and so its next order is five years later. (In the meantime, some of those 10 items have been counted as excess at the NADEP.) In that time, DLA has had no orders and, after four years, has declared as excess the remaining 90 units and disposed of them. Given this history, DLA considers the item a slow mover and must initiate a special purchase, and the feedback loop continues. Similar problems can arise for faster-moving items if the NSC maintains a stock of the item and ordinarily fills NADEP-replenishment requisitions from its own stock. In such cases, the NSC can fill several NADEP requisitions without generating any activity at DLA at all, so again DLA will see the item as a slow mover.

Inventory Management Using Value, Not Cost Alone

The discussion in the following sections of the *value* of a part is intended to help mitigate this disconnect between the NADEPs and the ICPs. Although the *value* of a part will be developed from the NADEP perspective, this measure could also be used by any of the activities that support NADEP work or any other remanufacturing or repair activity to manage the parts required to accomplish scheduled and unscheduled repairs.

The DoD currently emphasizes unit costing, but knowing the cost of the piece part is not enough. If a multiplier amplifier is keeping an F-14 from flying a sortie, its value may be considerably greater than its unit price of $7,470. If the part is never needed, its value is zero. A measure of value can help the Navy make appropriate decisions about what parts to stock. The value measure developed in the remainder of this report combines cost and actual system performance to yield a gauge that quantifies for a given investment the potential payoff in improved system performance.

3. METHODS FOR MANAGING REPAIR PARTS: THEORETICAL BASIS

What is an investment in the next unit of a repair part worth? The basis of our approach to the long-run problem is attributing a value to each unit of each repair part. But whereas repair parts are often valued by their unit price, unit price does not measure a part's worth: A buggy-whip has no value to the Navy, whatever its cost. We define the *value of a unit of a part* in terms of the benefit the Navy derives from it, where *benefit* is measured in days that end-items do not wait because of that part. (We will be more specific below.) Once value is expressed in days of end-item awaiting parts (AWP) time, it can be turned into familiar measures, such as "dollar value of the repair pipeline," for thinking about return on investments in parts.

This section discusses the theory underlying the definition and computation of value. First, we mention ways to use the idea of the value of parts. Next, we define value. Then, we show how to use value to efficiently stock parts at depots. Finally, we discuss details of how to compute value.

A FEW WAYS TO USE THE CONCEPT OF VALUE OF PARTS

The concept of value of parts has several applications. We briefly describe three such uses here. First, value can be applied to stock parts at depots or at intermediate repair facilities. Suppose a depot is going to invest a given amount in stocks of repair parts. For a given investment in inventory, an efficient investment stocks parts to obtain the most value (i.e., to keep jobs waiting as little as possible). Below, we discuss a method for making this investment, and in Section 5 we give an example using Navy data.

A second use of value is as follows. If repair parts can be valued in terms of days of AWP time and thus in terms of repair pipeline value, then investments in repair parts can be compared with other investments that reduce the value of the repair pipeline, such as retrograde transportation. A given dollar amount produces so much reduction if allocated to each of the possible investments that reduce the pipeline. An efficient investment allocates funds to produce the cheapest pipeline for the amount invested.

As a third use of value, suppose that a particular WRA, say, the radar antenna, is a problem. The question to be answered is, Which repair parts contribute most to the antenna's problem, and how can they be made less problematic? The method to be discussed below can be used to compute the value of the next unit of each repair part on the antenna. Sorting the parts in decreasing order of value is equivalent to sorting them in decreasing order of contribution to the WRA's problem. Suppose a circuit card is the part with the highest value. If its OST is reduced (see Section 6), its value decreases by some amount; if its replacement factor is reduced, its value decreases by a different amount. The relative payoffs of these or other fixes for the circuit card can be assessed by comparing the changes in value they induce.

Value of parts has other uses, but these three convey the possibilities. In the rest of this report we focus primarily on the first application: using value to stock parts.

DEFINING THE VALUE OF PARTS

Before we can define the value of parts, we need to describe the schematic of the depot that underlies the definition. Consider a WRA, say, the radar antenna, repaired by a particular depot. The process for fixing antennas is represented in Figure 1. Units of the antenna are inducted onto an antenna repair line. Repairing them involves removing NRFI

repairable subcomponents, SRAs, and replacing them with RFI SRAs, and removing and discarding NRFI piece parts and replacing them with RFI piece parts. Some types of SRAs are repaired at this depot. RFI units of those SRAs are obtained by inducting NRFI units onto an SRA repair line and repairing them, using piece parts and other SRAs. (Depots do not always have distinct WRA and SRA repair lines, but it is a convenient mental image.) Piece parts are either on hand at the depot or are obtained by placing requisitions to the supply system, which includes Naval Supply Centers (NSCs), supply depots of other services, Defense Logistics Agency (DLA) depots, commercial vendors, and local manufacture.

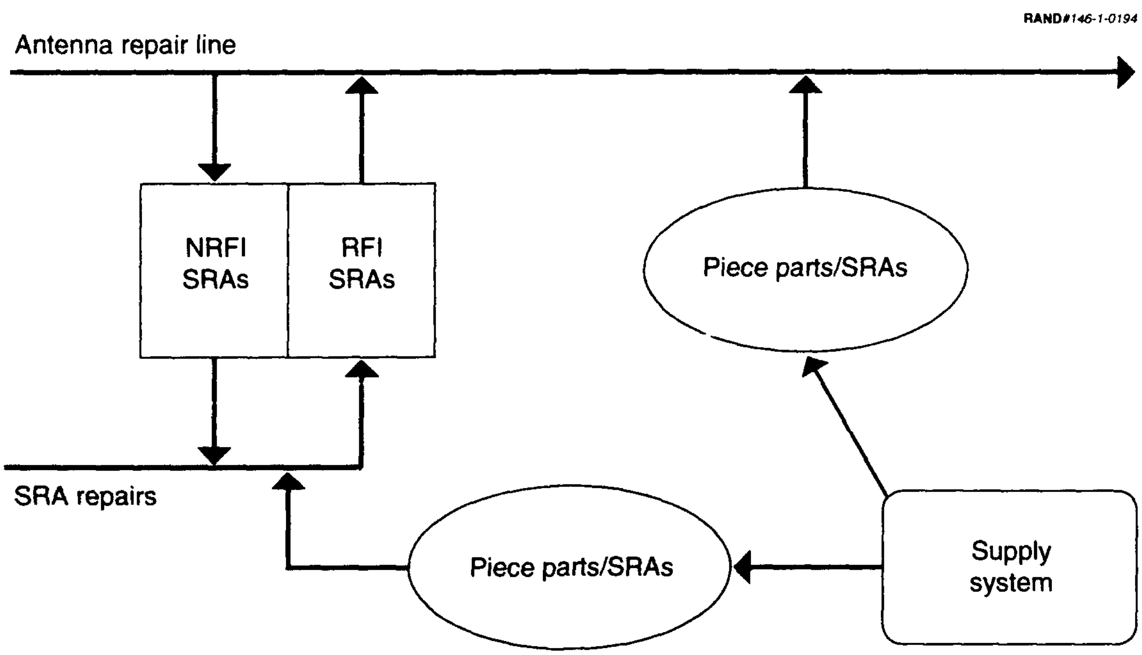

Figure 1—Schematic of Depot Repair

Repair parts flow out of depot stores onto antennas, and flow into depot stores in response to requisitions. The purpose of an inventory of parts—of authorized stock—is to put a buffer between the outflow onto antennas and the inflow from the supply system, so that repair jobs are inconvenienced as little as possible by requisition servicing or by repair of SRAs. If the buffer of authorized stock is made larger or more efficient, then the AWP time of WRAs will be smaller, so fewer WRAs will be in the repair pipeline at a given time. In other words, putting more or "smarter" parts into authorized stock takes WRAs out of the repair pipeline.

To see how stock affects the number of WRAs in the repair pipeline, consider Figures 2 and 3, depicting a notional WRA that uses a single repair part. Figure 2 shows the situation when authorized stock of that part is zero. Each time the WRA needs a unit of the repair part, it waits an OST for it. In the figure, OST is 40 days and the repair part is needed by some unit of the WRA every 30 days. One unit of the WRA is in AWP status two-thirds of the time and two units are in AWP status one-third of the time, so the average number of units in AWP status is 4/3. Figure 3 depicts the situation when the authorized stock of the repair part is one unit. The first WRA is serviced immediately, and a part is ordered to replenish the authorized stock. Although OST is still 40 days and a unit of the WRA needs the repair part every 30 days, each WRA (except the first) waits only 10 days instead of 40.

For 20 days out of every 30, no units of the WRA are in AWP status, so the average number of units of the WRA in AWP status is 1/3. Stocking one unit of the repair part removes one unit of the WRA from the repair pipeline.[13]

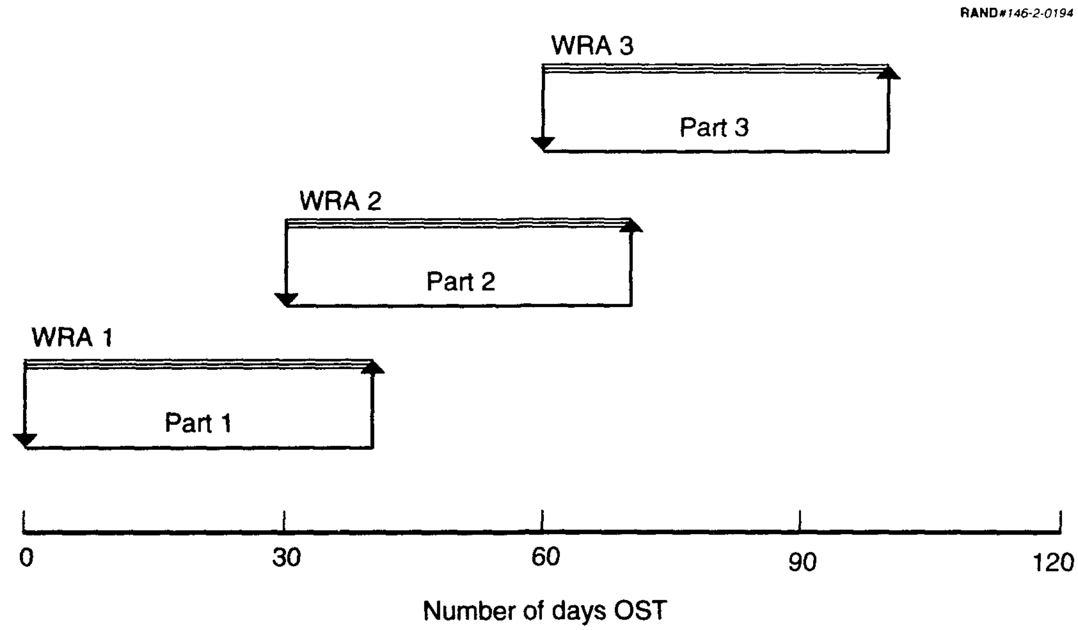

Figure 2—When Authorized Stock of the Repair Part Is Zero, Each WRA Waits One OST Each Time It Needs That Part

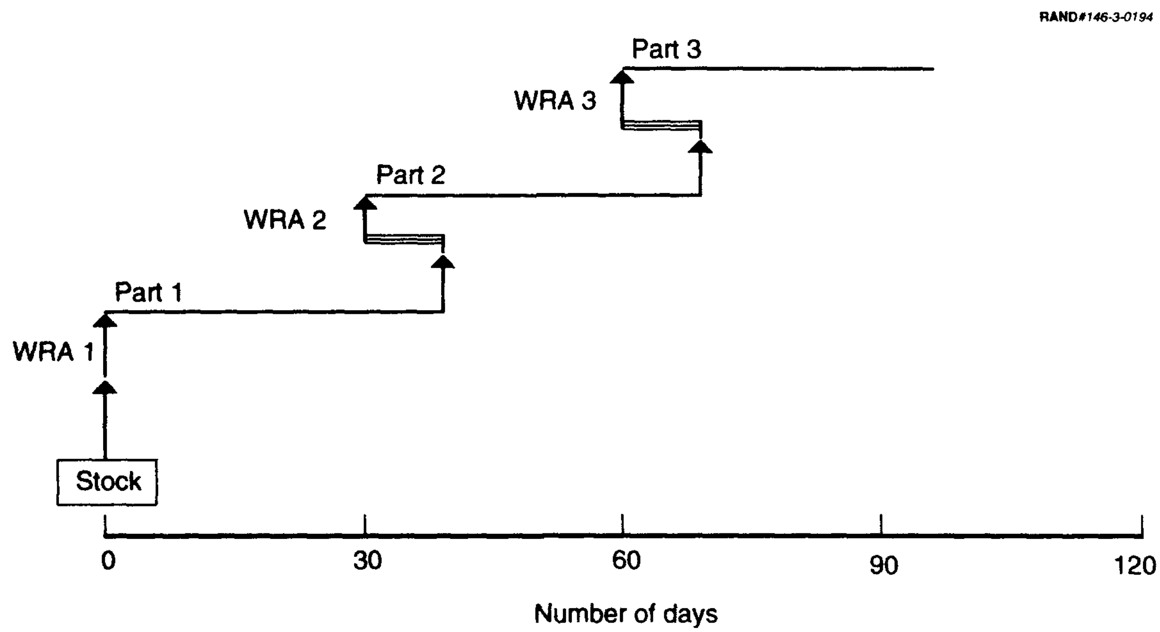

Figure 3—When Authorized Stock of the Repair Part Is One, WRAs Wait a Shorter Interval Each Time They Need That Part, and Fewer WRAs Are in the Repair Pipeline

We can now define *value*. A given set of stock levels yields particular expected AWP times for each WRA and a corresponding value for the repair pipeline. Add a unit of some repair part to authorized stock; the expected AWP time for the affected WRA will decrease, with a corresponding decrease in the value of the repair pipeline. The decrease in the value of the repair pipeline is the value of the unit added to authorized stock.

The foregoing discussion can be expressed mathematically. The standard formula for the value of the repair pipeline for the radar antenna is

$$\begin{aligned}&\{\text{unit price of the antenna}\} \times \{\text{no. of antennas in the repair pipeline}\}\\ &\quad = \{\text{unit price of the antenna}\} \times \{\text{antenna inductions per day}\}\\ &\quad\quad \times \{\text{days of turnaround time for the antenna}\}. \end{aligned} \tag{3.1}$$

To illustrate this last equality, suppose an NRFI antenna is inducted each day and that two days elapse between induction of an antenna and completion of repairs. (Suppose also that two antennas can be repaired at a time.) When the depot opens its antenna repair line, it has zero antennas in work. On the first day, an NRFI antenna is inducted into repair; therefore, the depot has one antenna in work, and after the second day, it has two antennas in work. At the beginning of the third day, an RFI antenna leaves repair, but an NRFI antenna is inducted. From then on, the depot always has two antennas in work. That is, its repair pipeline is two antennas, which is the product of the rate of antenna inductions (one per day) and the time each antenna spends in work (two days).

We assume that

$$\begin{aligned}&\{\text{days of turnaround time for the antenna}\}\\ &\quad = \{\text{days of AWP time for the antenna}\}\\ &\quad + \{\text{other components of turnaround time for the antenna}\}, \end{aligned} \tag{3.2}$$

and that stocks of parts affect only the AWP component of turnaround time. Among other things, this means that the processing time for an antenna is not affected by the list of parts that are needed to repair it, nor by their availability. Thus, stocks of parts affect only the contribution of AWP time to the value of the antenna repair pipeline, which is

$$\begin{aligned}&\{\text{unit price of the antenna}\} \times \{\text{antenna inductions per day}\}\\ &\quad \times \{\text{days of AWP time for the antenna}\}. \end{aligned} \tag{3.3}$$

If this argument is applied to all WRAs, then stocks of parts affect only the contribution of AWP time to the value of the depot repair pipeline, which is Eq. (3.3) summed across all WRAs. We actually compute the *expected* contribution of AWP time to the repair pipeline:

$$\begin{aligned}&\text{Sum}\ (\ \{\text{unit price of the WRA}\} \times \text{E}\{\text{inductions of that WRA per day}\}\\ &\quad \times \text{E}\{\text{days of AWP time for that WRA}\}\ ), \end{aligned} \tag{3.4}$$

where "E(X)" means "the mathematical expectation of X" and the sum is across all WRAs.

In these terms, the value of a repair part is defined as follows. For given part and antenna characteristics (to be discussed below) and given authorized stockage levels, the value of the next unit of a part used to fix the antenna is the reduction in the value of the antenna's repair pipeline from adding that unit to authorized stock:

$$\{\text{unit price of the antenna}\} \times E\{\text{antenna inductions per day}\} \times \Delta E\{\text{days of AWP time for the antenna}\}, \qquad (3.5)$$

where $\Delta E\{\text{days of AWP time for the antenna}\}$ is the change in the antenna's expected AWP time caused by the addition to authorized stock.

The value of the next unit of a part depends on the authorized levels of the other repair parts. Adding one unit of, say, a circuit card will cause different reductions in average AWP time, depending on the authorized stockage levels of all repair parts, including the circuit card itself. We discuss this more explicitly in the last subsection of this section—"Computing the Value of Parts."

Implicit in the notion of value discussed here is that the probability distribution describing part demands does not change over a long time. In reality, the values of parts are not static: They change over time as induction rates of WRAs change, as OSTs change, and so on. Thus, the efficiency of a set of authorized stockage levels also changes over time, so that an inventory that is efficient now will not be efficient in five years. (An *efficient inventory* is the set of authorized stockage levels that makes AWP time [actually, the value of the repair pipeline] shortest [smallest] for a given investment.) This change need not reflect badly on the people who buy the inventory. On the contrary, a successful program to improve reliability and maintainability will reduce the values of parts that are made less problematic. Depots must periodically reevaluate their authorized stock and consider changing it. This reevaluation is an important and complicated task that we do not have space to discuss here.

It is possible to define a notion of value more suitable to managing parts in the short run, for example, to solve the problem of how to allocate time and money for expediting parts in the next two weeks. We do not discuss this short-run problem here. It is discussed in detail in James S. Hodges, *Models and Algorithms for Repair Parts Investment and Management*, Santa Monica, Calif.: RAND, MR-314-A/USN, 1993; hereafter it is referred to as MR314.

STOCKING PARTS TO MAXIMIZE VALUE FOR A GIVEN INVESTMENT

Using this definition of the *value of the next unit of a part*, we can determine an efficient inventory. An algorithm to stock parts efficiently for a given investment can be specified simply, as follows:

- *Step 1.* Begin with zero authorized stock for all parts.
- *Step 2.* For each part, compute the value of the next unit of that part (we discuss how to do this in the next subsection) and divide its value by its unit price to obtain the rate of return from stocking that unit.
- *Step 3.* Select the part with the highest rate of return, and stock one unit of it.
- *Step 4.* Is the total cost of authorized stock less than the amount of money available to invest in parts? If yes, return to Step 2; otherwise, stop.

Steps 2 and 3 refer to evaluating and stocking the next unit of each item. It is not necessary to work with single units: Any increment to stock (the quantity usually purchased) can be evaluated. For cheap items, such as washers, or for items with units per application (UPA) greater than one, it will usually make sense to work with quantities greater than one unit. Step 4 discusses stopping the algorithm when a budget constraint is reached, but other criteria can be used to stop the algorithm. For example, stocking could continue until the

average AWP time is acceptable or until the return reaches some predetermined value, such as a 2:1 return.

The stockage problem we have posed is an instance of the classic knapsack problem.[14] Our algorithm is a heuristic solution, a so-called greedy algorithm, which has been extensively studied.[15] This greedy algorithm does not, in general, maximize the objective function subject to the constraints, but it is known to produce good solutions as long as the costs of the individual increments to stock are not large relative to the budget constraint. This optimization problem could be formulated as an infinite-stage dynamic program with an average reward criterion.[16] For simplicity, we did not use this formulation.

What Is the Relevant Measure of Cost?

Step 2 of the algorithm uses the unit price of each repair part as the measure of cost. To verify that this is the relevant measure, consider the following. We presume that the depot must buy every part it uses, including carcasses of SRAs. If the depot had authorized stockage levels of zero for all parts, it would need to buy from the supply system every part required for every repair job. If, instead, the depot has some authorized stock, it still must make the same sequence of buys from the supply system to replenish its authorized stock. The only difference is that the depot made a one-time purchase of parts up to authorized stockage levels, so that the repair lines receive parts more quickly. The purchase up to authorized stockage levels is an outlay made in addition to the cost of parts needed for jobs; it would be incurred whatever the authorized stockage levels were. The cost of authorized stock is determined by the depth of stockage[17] and by unit prices; the return for buying that inventory is reduced AWP time.

Implicit in this notion of authorized stock is that it will be replenished until the parts are removed from the inventory, making all the authorized stock obsolete. In reality, it will have some value as excess, so our measure of cost errs on the high side. However, if the scrap value of authorized stock is placed at some percentage of unit price, with the same percentage for all items, our measure of cost produces the right ranking of additions to authorized stock. By a similar argument, we do not explicitly account for holding costs.

Our notion of costs is most sensible for weapon systems that are just being moved to in-house depot repair, when the outlay for repair parts trades off directly for an outlay on WRAs to fill the repair pipeline. Our notion of costs and payoff cannot be taken literally for mature systems, because shortening the depot repair pipeline for such systems means that more RFI WRAs are available, not that fewer WRAs are procured to fill the repair pipeline. Computing actual savings for mature systems is complicated and is beyond the scope of this report.

COMPUTING THE VALUE OF PARTS

So far we have defined value and have specified an algorithm for stocking to maximize value without giving any details about how to compute value. This subsection discusses how we have implemented the value measure.

Some Specific Assumptions Needed to Compute Value

The ideas underlying the computation of value are illustrated by the hypothetical example in Figures 4–6. Figure 4 shows how requisitions are serviced if the depot has no authorized stock. Demands for a part, say, a circuit card for the radar antenna, occur on day 0 and on every third day afterward. If authorized stock is zero, then, each time a demand for the card occurs, a unit is requisitioned from the supply system and arrives one OST later,

which in the Figures is 7 days. The double line represents the amount of time each demand waits: 7 days.

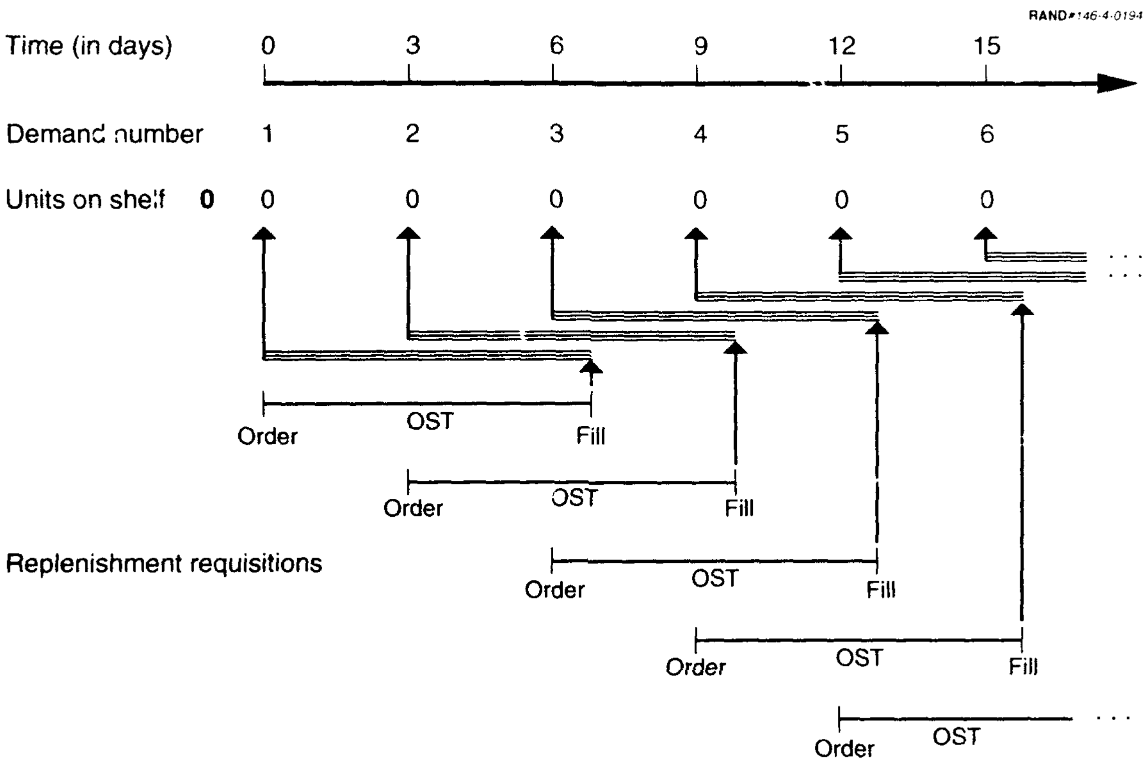

Figure 4—How Requisitions Are Serviced with No Authorized Stock

Figure 5 shows how requisitions are serviced if the depot is authorized to have one unit of stock of the card. At day 0, the depot has one unit on the shelf, which is used to fill the demand that occurs on day 0. A requisition is placed to replenish authorized stock; it is filled one OST later, on day 7. In the meantime, the second demand for the card occurs on day 3. This second demand is filled by the card that was ordered on day 0. After the second demand, a requisition is placed to replenish authorized stock; it is filled one OST later, on day 10. In the meantime, the third demand for the card occurs on day 6. And so on. As in Figure 4, the double line represents the amount of time each demand waits: With one unit of authorized stock, the wait is 4 days instead of 7.

Figure 6 shows how requisitions are serviced if the depot is authorized to have two units of stock of the card. At time 0, the depot has two units on the shelf. One of these units is used to fill the demand that occurs on day 0. A requisition is placed to replenish authorized stock; it is filled one OST later, on day 7. In the meantime, the second demand for the card occurs on day 3 and is filled by the second unit that was on the shelf at time 0. Another requisition is placed to replenish authorized stock; it is filled one OST later, on day 10. In the meantime, the third demand for the card occurs on day 6 and is filled by the unit that was ordered on day 0, so the third demand is kept waiting for one day. And so on. With two units of authorized stock, each demand waits one day.

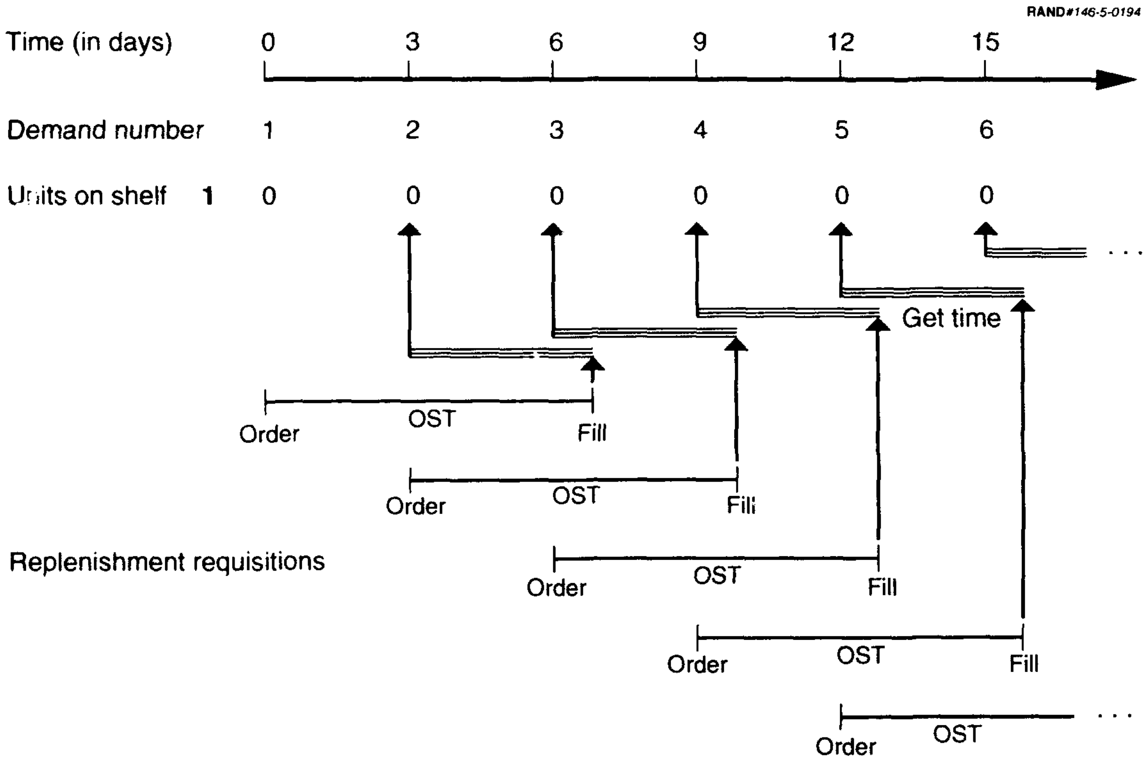

Figure 5—How Requisitions Are Serviced with One Unit of Authorized Stock

This example indicates the assumptions needed to compute the value of parts. We need to describe when, in the course of an antenna repair job, the need for the circuit card becomes known, and when in the repair the card is needed. These two times are needed to define AWP time. In the example shown in Figures 4 through 6, we finessed these data requirements by representing the occurrence of demands for the circuit card. To use our method, we need to describe WRA inductions and how part demands arise from those inductions. Although in Figures 4 through 6, each demand for the card triggered a replenishment requisition, we also need to know how depots replenish authorized stock. Finally, we need to describe OST. We have used particular assumptions in each of these areas, which we discuss in the paragraphs that follow.

When Part Requirements for a Job Become Known, and When in the Repair Those Parts Are Needed. Computing AWP time requires assumptions about the timing of repair-parts demands during repair jobs. Neither of our assumptions—that as soon as a WRA is inducted, the parts needed to repair it are known, and that the parts are needed immediately—is an accurate description of actual repairs. And they yield a definition of AWP time—that the clock measuring AWP time starts as soon as a WRA is inducted and runs until the last repair part arrives—that is not consistent with Navy terminology. They will be least accurate for end-items for which processing is complicated and for which parts are not needed until late in the process, e.g., engine rework. Neither of our assumptions is necessary to compute value, but they greatly simplify the computation and substantially reduce its data requirements. In particular, unless the sequence and timing of actions within a job are understood and measured well, we believe it is pointless to try to replace our assumptions with assumptions having more face validity.

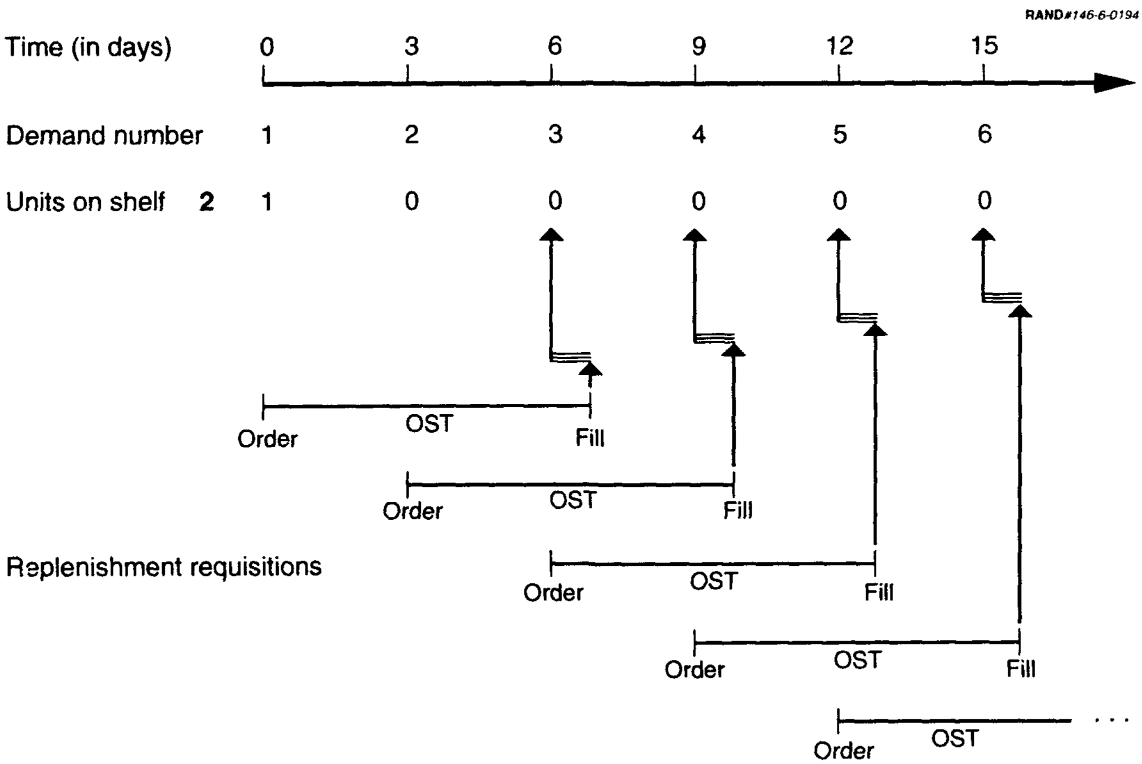

Figure 6—How Requisitions Are Serviced with Two Units of Authorized Stock

How Depots Replenish Authorized Stock. We assume that a replenishment requisition for part i is placed after every sth demand for part i. This is often called (S – s, S)[18] ordering, where S is the authorized stockage level. If S = 0—the part has no authorized stock—then s is set to 1 and an order is placed after each demand. As Figures 4 through 6 suggest, given the long OSTs currently experienced by the Navy, a depot usually will not have the authorized stock of a part on the shelf, particularly for high-demand items. Instead, units are demanded for repair jobs and replacement units are ordered; if demands occur in an orderly fashion, the bulk of authorized stock will be in the due-in pipeline. Thus, the buffer of authorized stock serves two functions: (1) Stock on the shelf is insurance against demands for slow-moving items and (2) it is insurance against surges in demand for faster-moving items. Stock in the due-in pipeline has, in effect, been proactively ordered. If OSTs were radically shortened, most of the authorized stock would be on the shelf at any time.

Induction of WRA Jobs and OST. In MR314, we discuss three sets of assumptions about induction of WRA jobs, OST, and SRA repair. The sets vary in the amount of uncertainty they incorporate. Here is the set of assumptions used in this document (called the "U" assumptions in MR314):[19]

- Induction of units of WRA k is a homogeneous Poisson process with rate λ_k.
- OST of part j ordered from the supply system follows a negative exponential distribution with mean $1/\delta_j$.

The other sets of assumptions incorporate less uncertainty about the timing of WRA inductions and OST. We make several additional assumptions, which we discuss at the end of this section.

Computing the Value of a Unit of an Item

Consider the circuit card going on the radar antenna. We want to assess the value of the next unit of that circuit card.

For given characteristics,

- Part characteristics: RFs, OSTs, unit prices, and (for SRAs) bills of materials
- Antenna characteristics: bill of materials, unit price, and daily induction rate
- Authorized stockage levels for all parts,

the value of the next unit of the circuit card is computed by the following four steps:

1. For each SRA that is repaired locally, compute the expected AWP time for that SRA when it is repaired:
   — For each part going on the SRA, compute the expected time the SRA waits for that part if it needs one (using formulas discussed below). This expected waiting time depends on the number of units of that part in authorized stock.
   -- Plug the expected time and the replacement factors (RFs) of the parts into the tall-pole formula, discussed below.

2. Compute the expected AWP time for the antenna:
   — For each piece part going directly on the antenna and each SRA not repaired locally, compute the expected wait for that part if the antenna needs one (using formulas discussed below). This expected waiting time for each part depends on the number of units of that part in authorized stock.
   — For each SRA repaired locally, compute the expected wait for that SRA if the antenna needs one, using the expected AWP time computed in Step 1 and formulas discussed below. This expected waiting time depends on the number of units of that SRA in authorized stock.
   — Plug these times and the RFs of the parts that go on the WRA into the tall-pole formula.

3. Increment the circuit card's authorized stock by one unit, repeat the first two steps, and compute the change in the antenna's expected AWP time due to the extra card.

4. Compute the value of the next unit of the circuit card by multiplying the change in expected AWP time, obtained in Step 3, by the antenna's daily demand rate and unit price.

The computation just described assumes that only piece parts are used to repair SRAs, which is not true in general. Our method can be extended to more levels of indenture; however, as discussed in Appendix B of MR314, extending the computation to more levels of indenture may not be advisable.

The description just given assumes that we have two formulas: (1) The "tall-pole formula" for computing expected AWP time given RFs and expected waits for individual parts, and (2) a formula for the expected wait for a part if that part is needed. The second formula is described in Section 4 and Appendix A of MR314. The remainder of this section describes the tall-pole formula and gives the expected wait for a part if it is needed and the way that characteristics of parts and WRAs determine the value of parts.[20]

Table 1

Part Characteristics for a Hypothetical WRA

NIIN[a]	RF	Wait If Needed
0001	0.05	40
0002	0.10	35
0003	0.50	25
0004	0.01	10

[a]National item identification number.

The Tall-Pole Formula

We use the tall-pole formula repeatedly in computing value. To understand it, consider a hypothetical WRA with four piece parts, described by Table 1. For each of the four parts, we have the RF and the amount of time that the artisan will wait for the part if it is needed. For now, assume that these waiting times are deterministic. If the repair job requires NIIN 0001, then no matter what other parts are demanded, the job will wait 40 days for NIIN 0001; this part will be the "tall pole in the tent." If the repair job does not require NIIN 0001 but does require NIIN 0002, then regardless of whether NIINs 0003 and 0004 are demanded, NIIN 0002 is the tall pole in the tent and the job waits 35 days. If neither NIIN 0001 nor NIIN 0002 is required, but NIIN 0003 is, then it is the tall pole and the job waits 25 days. If only NIIN 0004 is demanded, the job waits 10 days. Therefore, the mathematical expectation of the time the job waits for parts is

$$
\begin{aligned}
\text{E(wait)} &= \text{Sum(\{wait if part i is the tall pole\}} \\
&\quad \times \text{Probability\{part i is the tall pole\})} \\
&= 40 \times \text{Pr(need NIIN 0001)} \\
&\quad + 35 \times \text{Pr(need NIIN 0002, do not need 0001)} \\
&\quad + 25 \times \text{Pr(need NIIN 0003, do not need 0001 or 0002)} \\
&\quad + 10 \times \text{Pr(need NIIN 0004, do not need 0001, 0002, or 0003)} \\
&\quad + 0 \times \text{Pr(do not need any parts)} \\
&= 40 \times .05 + 35 \times .10 \times .95 + 25 \times .50 \times .95 \times .90 \\
&\quad + 10 \times .01 \times .95 \times .90 \times .50 \\
&= 16.06 \text{ days.}
\end{aligned}
\tag{3.6}
$$

This is the tall-pole formula for the expected amount of time a WRA (or SRA) waits for repair parts. If the waiting times for individual parts are deterministic, the tall-pole formula gives the exact expected value. If the waiting times for individual parts are stochastic—so that the column "wait if needed" in Table 1 gives the *expected* wait if needed—then the stochastic (random) process yielding the WRA's wait for parts is much more complicated and the tall-pole formula is approximate. This complication arises because, first, the parts required for the repair are determined stochastically, then waiting times for the needed parts are determined stochastically, then the WRA's waiting time is determined as the maximum of those waiting times. As the number of parts on the WRA grows, it becomes much more time-consuming to compute the expectation of this maximum. The tall-pole formula approximates the exact computation. The quality and consequences of this approximation are discussed in Appendix B of MR314.

How WRA and Part Characteristics Determine the Value of Parts

Consider a hypothetical radar antenna composed of two parts, a circuit card and a bolt. Recall the definition of the value of the next unit of the circuit card:

$$\{\text{unit price of the antenna}\} \times \{\text{antenna inductions per day}\} \\ \times \Delta E\{\text{days of AWP time for the antenna}\}, \tag{3.7}$$

where ΔE(days of AWP time for the antenna) is the reduction in the antenna's expected AWP from adding a unit of the circuit card to authorized stock. The card's value depends on the following factors:

- The unit price of the antenna
- The induction rate of the antenna
- The effect of the characteristics of the circuit card and the bolt on ΔE{days of AWP time for the antenna}.

Value is proportional to the unit price of the antenna. Suppose the circuit card and another part, say, a diode, could each eliminate a day from the AWP time of, respectively, the antenna and a converter-programmer. If the antenna and converter-programmer fail at the same rate, then the Navy benefits more by saving the day on the more expensive WRA.

The induction rate of the antenna and the characteristics of the circuit card and the bolt have more complicated effects on the value of the next unit of the card. To understand these effects, consider the tall-pole formula for our hypothetical radar antenna. If the expected wait for the bolt is longer, the tall-pole formula yields

$$\begin{aligned} E(\text{antenna AWP}) &= RF(\text{bolt}) \times E(\text{wait for bolt if needed}) \\ &+ (1 - RF(\text{bolt})) \times RF(\text{card}) \times E(\text{wait for card if needed}). \end{aligned} \tag{3.8}$$

If the expected wait for the card is longer, the tall-pole formula yields

$$\begin{aligned} E(\text{antenna AWP}) &= RF(\text{card}) \times E(\text{wait for card if needed}) \\ &+ (1 - RF(\text{card})) \times RF(\text{bolt}) \times E(\text{wait for bolt if needed}). \end{aligned} \tag{3.9}$$

Thus, ΔE(antenna AWP) is expressed by one of the following two formulas:

$$RF(\text{card}) \times \Delta\, E(\text{wait for card if needed}), \tag{3.10}$$

if the wait for the card is longer.

$$(1 - RF(\text{bolt})) \times RF(\text{card}) \times \Delta\, E(\text{wait for card if needed}), \tag{3.11}$$

if the wait for the bolt is longer.

These two formulas show how the value of the card depends on the bolt's characteristics. If the expected wait for the card (given its authorized stockage level) is longer than the expected wait for the bolt (given its authorized stockage level), the value of the next unit of the card does not depend explicitly on the bolt's characteristics. If the expected wait for the bolt is longer than the expected wait for the card, the value of the card is determined by the bolt's RF. For example, if the expected wait for the bolt is 30 days and the expected wait for the card is 25 days, *and* antennas always need the bolt—that is, RF(bolt) = 1—then the next unit of the card has zero value, because the antenna waits longer for the bolt than for the card regardless of whether another card is stocked. If only half the antenna jobs need the

bolt—RF(bolt) = 0.5—then the AWP times of the other half of the antenna jobs can potentially be reduced by reducing the waiting time for the circuit card; therefore, the next unit of the circuit card has some value.

The effects of the induction rate of the antenna and the circuit card's characteristics are somewhat more complicated. From Eq. (3.10), the value of the next card is proportional to

$$\{\text{induction rate for antenna}\} \times \text{RF(card)} \times \Delta\text{E(wait for card if needed)}. \tag{3.12}$$

(Actually, Eq. [3.12] is true only as long as the expected wait for the card and the expected wait for the bolt remain in the same order; when the stock of the card increases to the point at which its expected wait becomes smaller than the bolt's expected wait, the card's value changes discontinuously.) The third factor in this expression, ΔE(wait for card if needed), is determined by the assumptions made about antenna inductions and OST. For the "U" assumptions made in MR314,

E(wait for card if authorized stock = n) =

$$E\{OST(card)\}\left(\frac{\text{RF(card)} \times \text{I(antenna)}}{\text{RF(card)} \times \text{I(antenna)} + 1/E\{OST(card)\}}\right)^{n}, \tag{3.13}$$

where I(antenna) is the daily induction rate of the antenna and E{OST(card)} is the expected OST of the card. Therefore,

$$\Delta\text{E(wait for card if needed)}$$

$$= \text{E(wait for card if authorized stock} = n) - \text{E(wait for card if authorized stock} = n + 1)$$

$$= \frac{(\text{RF(card)} \times \text{I(antenna)})^{n}}{(\text{RF(card)} \times \text{I(antenna)} + 1/E\{OST(card)\})^{n+1}}. \tag{3.14}$$

Then the value of the next card is proportional to

$$\left(\frac{\text{RF(card)} \times \text{I(antenna)}}{\text{RF(card)} \times \text{I(antenna)} + 1/E\{OST(card)\}}\right)^{n+1}. \tag{3.15}$$

As the authorized stock of the card increases, the value of the next unit decreases—that is, the marginal return to stocking the card is decreasing. As the expected OST of the card increases, the value of the next unit of the card increases. The derivative of value with respect to RF is positive, so the card's value increases with its RF. (The second derivative is negative, so that the increase in value with RF slows as the RF becomes larger.) Symmetrically, as the antenna's induction rate increases, so does its value. (The second derivative of value with respect to induction rate is also negative.) If the antenna never breaks—I(antenna) = 0—or if the antenna breaks but the circuit card never does—RF(card) = 0—there can be no value to stocking the circuit card. If either the antenna's induction rate or the card's RF is increased, an antenna job will wait longer each time a circuit card is demanded, so the value of the next unit of the card is increased.

Miscellaneous Other Assumptions

Our method uses other assumptions that may be important. We assume the following:

- The depot has adequate capacity to repair the WRAs that are scheduled.
- The number of units of a WRA in work is small relative to the total number of units owned by the Navy. The effect of this assumption is that the induction rate of WRAs is not affected by the number in the repair pipeline.
- Due-in parts are not used opportunistically; that is, we assume FIFO. In Figure 4, this means that the part ordered for demand 1 is used to fill demand 1, even if using it to fill demand 2 would make a WRA available sooner.
- SRA repair time = process time + AWP time, where process time is unaffected by the availability of parts.
- No cannibalization occurs.
- No condemnation of WRAs or SRAs occurs.
- UPA is 1.
- There are no common items.[21]

A complete list of assumptions of the long-run methods is given in MR314, Appendix B. As noted above, each individual part follows a queueing process, with the specific process depending on the specific probabilistic forms for WRA inductions, OST, and the occurrence of part demands for individual WRA or SRA jobs. The new wrinkle in our setup is that, for each WRA or SRA repair job, the AWP time is the maximum across several queues selected stochastically from a larger group of queues, corresponding to the parts of which the WRA or SRA is composed.

4. USING NAVY DATA TO BUILD AUTHORIZED STOCKS AT DEPOTS

The preceding sections outline our goals for this project and the mathematical theory underlying our method for managing repair parts. The rest of this report identifies the steps necessary to implement the method and presents some illustrative calculations. This section discusses the steps required to build authorized stocks using our method; Section 5 discusses how we evaluated stocks constructed using our method.

The output of our method of constructing authorized stocks is a list of increments to authorized stock, sorted in order of desirability, as defined in Section 3. Three steps are required to build such a list:

1. Compile selected data about WRAs, SRAs, and piece parts.
2. Build a rank-ordered list of increments to authorized stock.
3. Cut off the list to determine the stock to be authorized.

Each of these steps is described in a subsection to follow.

COMPILE DATA ABOUT WRAs, SRAs, AND PIECE PARTS

The main data source available to us was the Naval Inventory Materiel Management System (NIMMS), which records transactions for materiel at all the NADEPs. NIMMS was our sole source of information on subcomponents of WRAs. We used miscellaneous ASO files for information about weapon systems and about the WRAs themselves. We used three years of NIMMS (1989–1991) from the Jacksonville, Norfolk, and North Island NADEPs. The ASO data we used are largely insensitive to time; we accessed those data items in approximately August 1991.

We summarize supply system performance by weapon system—actually, by special materiel identification code (SMIC), which corresponds roughly to major systems. We selected 11 SMICs on the basis of whether our NIMMS data contained many repair jobs (more than 1,000) for WRAs with those SMICs. Table 2 gives the 11 SMICs and the number of WRA repair jobs we used for each.

Table 2

SMICs Used in the Analysis and the Number of WRA Repair Jobs

SMIC	Description	Number of Jobs
BE	E2/C2 electronic aircraft	4,526
BP	P3 patrol aircraft	7,915
CY	AWG-9 radar	3,408
DH	H3 helicopter	1,556
DQ	T56 engine	6,048
EQ	T58 engine	9,055
FQ	T64 engine	3,591
MH	H46 helicopter	2,524
PQ	TF30 engine	6,515
RA	A6E attack aircraft	3,466
TN	F404 jet engine	5,158
TOTAL		53,762

Background on NIMMS

NIMMS logs every materiel transaction made by NADEP artisans and by NADEP supply organizations, such as the Naval Industrial Fund (NIF) store. The transactions are primarily parts requisitions and issues, but NIMMS is also intended to log other materiel actions, such as turn-ins of unused material and cannibalization of parts. (Not all these functions are used, however.) When a job is initiated at a NADEP, it is assigned a unique job-number/link-number combination. This job/link combination is recorded in NIMMS on all requisitions for parts for that job. The last four digits of the job number are the item identification code (IIC) of the item being repaired, so that job numbers can be used to identify the WRA or SRA being repaired. Job numbers also identify whether the job is component repair or standard depot-level maintenance (SDLM);[22] our subset of NIMMS data excluded the latter.

The history of a requisition is recorded in NIMMS as a series of transactions sharing a common document number. Each transaction contains the document number and the transaction date, so it is possible to build up a log, by job, of which parts are requested by artisans, when they are requested, and the time until the request is satisfied (if it is). A partial example of a common document is shown below. Requisitions of items that are not stocked in the NIF store usually have three transactions: the first records the request (a request for NIIN [national item identification number] 000084476 is made on 91224 [Julian date]; the code 11 record indicates that the item was not in stock); the second records receipt of the item by the NIF store (8 days after the request, NIIN 000084476 was received from the supply system, on 91232); and the third records issue of the item to the artisan (code 32). In the second example, a request for NIIN 001530123 was satisfied by the NIF store (only one record for NIIN 001530123 and DOCNO 110610B8 with code 32).

IIC	NIIN	DOCNO	DATE	CODE
H502	000084476	122413B8	91224	11
H502	000084476	122413B8	91232	31
H502	000084476	122413B8	91232	32
H502	001530123	110610B8	91106	32

Data Required by Our Stockage Method

To use the value of parts method (described in Section 3) to stock parts, we need the data listed in Table 3. This subsection documents how we constructed the data items in Table 3, providing a basis for discussing deficiencies in available Navy data sources that made it difficult to evaluate our stockage method, as is done in Section 5. As will be seen, NIMMS is an imperfect data source for some of the quantities in Table 3, but it is our only source. Some deficiencies in the data forced us to forgo certain features of the stockage method, which we indicate below.

(a) WRA: SMIC. For every job that appeared in our NIMMS data, we identified the corresponding WRA and assigned that WRA to a SMIC. NIMMS identifies WRAs by the IIC, and an ASO file called BX1 enabled us to map IIC codes approximately to weapon systems through the SMIC, as listed in Table 2. BX1 also enabled us to map various identifiers, such as NIINs, into IICs. If a WRA's IIC did not appear in the BX1 file, or if it did not correspond to one of the 11 SMICs in Table 2, we dropped the WRA from further consideration.

Table 3

Stockage Method Data Requirements

Level		Type of Information
WRA	(a)	SMIC of WRA, various WRA identifier mappings
	(b)	Rate of induction for repair
	(c)	Unit price
	(d)	Bill of materials (i.e., a list of SRAs and piece parts that go on the WRA, and their replacement factors)
SRAs repaired locally	(e)	Bill of materials (i.e., a list of piece parts that go on the SRA, and their replacement factors)
	(f)	Repair processing time (including AWP time)
	(g)	Unit price
Piece parts and SRAs not repaired locally	(h)	Unit price
	(i)	Order-and-ship time

(b) WRA: Rate of Induction for Repair. Each job in NIMMS has a unique job/link identifier. We counted the number of unique jobs over the three-year period for each end-item (i.e., WRA or SRA), divided that by 1,095 (365×3) to get a rate of inductions per day, and postulated the resulting value as the daily rate of inductions. It is not necessary to supply the WRA induction rates this way. Rates could be forecast by some other method: For example, a "tired iron" program would increase the rate of induction of affected WRAs, and those higher rates could be used in our stockage method in place of historical rates.

(c) WRAs: Unit Price. The ASO BX1 file provided this information directly in the data element STDPRICE.

(d) WRAs: Bill of Material. Until recently, reliable BOMs were available for few Navy components. With the recent introduction of the Automatic Bill of Materials System and NADEP Logistics Management System, BOMs have been constructed or are being constructed for most Navy components. However, our early attempts to use the output from this process were unsuccessful because identifiers between the constructed BOMs and NIMMS were inconsistent. We decided to construct our own BOMs for parts going on a WRA from NIMMS data, as follows: We observed in NIMMS which parts were ordered against each WRA, then estimated replacement factors as the number of orders for each part on a WRA divided by the number of jobs for that WRA.

This latter method is not perfect. Some end-items appear only a few times, so the estimates of their replacement factors are unreliable. Also, in three years, not all the parts on a WRA will be requisitioned. Moreover, NIMMS does not check that requisitioned parts actually go on the WRA against which they are charged, which reportedly allows some abuses, e.g., artisans sometimes cross-order parts to build up a bench stock in anticipation of supply failures. Consequently, parts in NIMMS can be charged against WRAs that they do not actually go on.

(e) SRAs: Bill of Materials. There are two kinds of SRAs: those that are repaired at NADEPs and those that are repaired by contractors. For our purposes, the latter are treated as piece parts. Theoretically, the former can be identified from the IIC in the job numbers in NIMMS. From conversations with NADEP personnel and our observations from the NIMMS data, it appears that few SRA repair jobs are recorded in NIMMS (at least during the period

we used). For example, concurrent repair of SRAs is done for aircraft in SDLM, and parts needed for such repairs are ordered in various ways that are difficult or impossible to trace.

The result is that we were unable to get transaction-level data for SRA repairs; therefore, we do not have data on which SRAs are repaired at the NADEPs or what their BOMs are. Lacking a data source, we were unable to represent SRA repairs in the sample calculations shown later in this section or in the evaluation discussed in Section 5. If the data are available from another data source, or become available from NIMMS, however, SRA repairs could be taken into account in using our stockage method.

(f) SRAs: Repair Processing Time. NIMMS provides no information on repair pipeline times. We approached the NADEPs for such information, but found that we were unable to get AWP times for specific SRAs (or for WRAs, for that matter). Attempts were being made to collect these data in certain shops, but the information was incomplete when we needed it. We judged that we could not come up with even a crude approximation to repair times for our computations. This lack of data also contributed to our decision not to represent SRA repair in the sample calculations in this section or in the evaluation in Section 5.

(g) SRAs: Unit Price. The ASO BX1 file provided the unit price for SRAs directly in the data element STDPRICE.

(h) Piece Parts and SRAs Not Repaired Locally: Unit Price. Each NIMMS record contains the unit price of the requisitioned part. A NIIN requisitioned n times will have n NIMMS records, each containing the price of that part. In some cases, the prices were not the same on all records. To smooth out this variation, we computed the median price across all transactions in NIMMS for each NIIN, then computed max(price, $1) to effectively bound the NIINs' rates of return.

(i) Piece Parts and SRAs Not Repaired Locally: Order-and-Ship Time. NIMMS records the date a requisition was placed, in the document number of that requisition. It also records a date for each transaction on that document. We computed the amount of time it took to service each requisition as the difference between the date in the document number and the date on the issue transaction. To estimate OSTs, we used only requisitions for items not stocked in the NIF store, i.e., requisitions that were not satisfied in the local NIF store or were not back-ordered NIF-stocked items.[23]

Some requisitions in NIMMS were apparently never filled: Their documents had no issue transaction. We could use special statistical estimation techniques to incorporate those requisitions in an estimate of average service time—if we believed that they were genuinely open at the end of the period covered by our data. However, consistent features of the NIMMS data lead us to believe that unfilled requisitions are usually the result of an artisan's not telling NIMMS that the part had been received or that the problem had been worked around. For example, some requisitions are open for years, whereas subsequent requisitions for the same part are filled in just a few days. We chose to ignore these cases rather than treat them at face value as requisitions with very long wait times.

We could have used a distinct OST for each NIIN. However, many NIINs had either few requisitions or considerable variation in service times, and we wanted to smooth out this variation. Therefore, we averaged OSTs within SMIC by ICP, as identified by the cognizance code (COG).[24] Table 4 contains the average OSTs we used. Where average service time was greater than 120 days, we substituted 120 days; where it was zero, we substituted 1 day. We did this to reduce the effect of outlying data values from small sample sizes on our computations.

Table 4

OSTs (days), by COG and SMIC

	SMIC										
COG	BE	BP	CY	DH	DQ	EQ	FQ	MH	PQ	RA	TN
0Z		26.6									
1H	16.1	48.0	84.5	15.3	7.2	17.0	22.5	1.0	8.0	23.0	
1R	48.5	24.9	22.6	39.5	15.0	21.4	32.9	23.4	12.3	20.4	35.2
5R		2.0					8.0				
6K	12.0	120.0	24.4	8.0					19.3	75.2	8.0
6V	13.8	30.0	9.7	30.0	22.7	24.5	12.9	21.8	15.0	29.3	[illegible]
7R	36.0	120.0	40.0	24.0	31.0	17.9	36.4	74.6	24.9	64.5	[illegible]
9A		6.0		11.8							
9C	33.5	18.5	19.9	14.8	13.7	20.2	23.0	16.1	15.4	21.2	20.5
9D		1.0		12.0		47.0		57.0	7.0	9.0	
9F	21.3	17.6	24.2	19.4	1.0			5.0	7.5	11.9	120.0
9G	21.8	22.9	12.9	23.0	19.3	27.0	15.3	21.6	22.7	17.5	41.5
9I	16.0	1.0		14.2			45.3	45.8	120.0	20.4	
9J	24.8	13.3	32.4	13.7	27.2	12.4	15.3	15.2	14.7	19.0	19.8
9K	23.4	10.0	69.4		17.5	19.7	5.5	11.0	4.1	14.6	14.7
9L	14.0	13.6	21.1	19.2	26.0	1.0	14.0	23.0	9.9	8.3	1.0
9N	20.4	18.6	13.4	20.5	13.3	10.9	18.8	18.4	13.4	12.4	22.2
9Q	14.6	20.0	17.4	33.1	22.1	42.0	1.0	23.2	14.9	18.4	29.0
9S	10.0	17.5								10.5	
9V	10.3	19.0	16.6	12.1	16.6	4.8	99.1	24.3	21.6	18.3	19.2
9W	13.1	15.5		24.9	13.3	18.0		15.1		14.0	22.8
9X	13.0	1.0							1.0		14.0
9Y		12.2	11.6	14.0	2.0	8.0			6.0	8.0	
9Z	25.5	18.7	25.0	26.0	16.2	17.2	23.0	19.9	15.2	18.4	22.9

NOTES: Blanks signify SMIC/COG combinations that did not occur in NIMMS. Mean values are truncated below at 1.0, above at 120.0.

Implications of Using NIMMS Data

The data available to us shaped the exact way in which we implemented our stockage method. If different data were available, we could implement the value of parts idea in a different way. For example, if data were available on processes rather than on jobs, we could use the value idea to reduce the time a WRA spends in a particular process or to identify processes that were the bottlenecks in NADEP throughput. As noted earlier, lack of data about SRA repair precluded modeling it as described in MR314. Additionally, the imperfections in NIMMS affect the kinds of conclusions we can draw from any computations based on NIMMS data. We discuss this issue at greater length in Section 5.

BUILD A RANK-ORDERED LIST OF INCREMENTS TO AUTHORIZED STOCK

Having compiled the data, the next step is to build a list of increments to authorized stock. As Section 3 indicates, we do this by setting each part's authorized stock to zero, then repeatedly selecting the item with the highest return and adding a unit of that item to the authorized stock. We go through an example to show how the computation works for an individual WRA; then, we show part of the sorted list obtained from doing this calculation for all the WRAs in each SMIC.

Table 5

Part Characteristics for a Hypothetical WRA

NIIN	RF	Get Time (days)	Unit Price
0001	0.25	31	$400
0002	0.10	20	4
0003	0.15	15	200
0004	0.07	10	50

Sample Computation for a Single WRA

Assume that a hypothetical WRA has four parts and is inducted at the rate of one every two weeks. We begin the computation with authorized stockage levels of zero for each of those parts. Thus, if a repair job needs a part, it must wait one OST before that part is available. Table 5 gives the parts' replacement factors, unit prices, and "get times," i.e., the amount of time an artisan must wait, on average, to obtain the part if he or she needs it. We make three assumptions: that demands for each of these four NIINs are independent; that these parts are ordered at the start of the job; and that the get times are nonstochastic.[25] Then the expected time that the WRA waits for parts (what we call the tall-pole computation, which is described in Section 3) is

$$\begin{aligned} E(AWP) = \ & 31 \times \Pr(\text{need NIIN 0001}) \\ & + 20 \times \Pr(\text{need NIIN 0002, do not need 0001}) \\ & + 15 \times \Pr(\text{need NIIN 0003, do not need 0001 or 0002}) \\ & + 10 \times \Pr(\text{need NIIN 0004, do not need 0001 or 0002 or 0003}) \\ = \ & 31 \times .25 + 20 \times .10 \times .75 + 15 \times .15 \times .75 \times .90 \\ & + 10 \times .07 \times .75 \times .90 \times .85 \\ = \ & 11.17 \text{ days.} \end{aligned} \tag{4.1}$$

Table 6

Part Characteristics for a Hypothetical WRA, with Tall-Pole Results Appended

NIIN	RF	Old Get Time	Unit Price	Old E(AWP)	New Get Time	New E(AWP)	E(AWP) Reduction	E(AWP) Reduction/$
0001	0.25	31	$400	11.17	11.04	6.54	4.63	0.012
0002	0.10	20	4	11.17	2.50	10.03	1.14	0.285
0003	0.15	15	200	11.17	2.08	9.92	1.25	0.006
0004	0.07	10	50	11.17	0.48	10.79	0.38	0.008

NOTE: Get time and E(AWP) are in days.

The rest of the computation amounts to a sequence of spreadsheet operations that add columns to Table 5, the result of which appears as Table 6. First, add a column containing the expected maximum wait, just computed (called "old E(AWP)"). Next, for each of the four items, we ask the following question: What would the expected wait be if we stocked one more unit of that item? (In our example, this would be the first unit of authorized stock for each item.) For concreteness, consider NIIN 0001. To determine the expected wait if we

stocked one unit of NIIN 0001, we need to know what the average get time for NIIN 0001 would be if the stockage level were one instead of zero. From Section 3, we know it is shorter than the original 31 days. From Eq. (3.14), we know the answer is 11.04 days, so that it is now the third-longest get time instead of the longest. Then the WRA's new expected wait is computed using the tall-pole computation:

$$\begin{aligned}\text{new } E(AWP) = {} & 20 \times \Pr(\text{need NIIN 0002}) \\ & + 15 \times \Pr(\text{need NIIN 0003, do not need 0002}) \\ & + 11.04 \times \Pr(\text{need NIIN 0001, do not need 0002 or 0003}) \\ & + 10 \times \Pr(\text{need NIIN 0004, do not need 0002 or 0003 or 0001}) \\ = {} & 20 \times .10 + 15 \times .15 \times .90 + 11.04 \times .25 \times .90 \times .85 \\ & + 10 \times .07 \times .90 \times .85 \times .75 \\ = {} & 6.54 \text{ days.} \end{aligned} \tag{4.2}$$

The foregoing computation is represented in Table 6. For NIIN 0001, the column "new get time" gives the amount of time a WRA job waits for NIIN 0001 if its authorized stock is one unit: 11.04 days. Similarly, for NIIN 0001, the column "new E(AWP)" gives the expected time the WRA waits for parts if the authorized stock of NIIN 0001 is one unit: 6.54 days. For each NIIN, the entries under these two new columns are, respectively, the get time for each part if the authorized stock is one unit, and the expected time the WRA waits for parts if the authorized stock is one unit.

The next column in Table 6 is "E(AWP) reduction," which is old E(AWP) minus new E(AWP), i.e., for NIIN 0001, it is the reduction in the expected time the WRA waits if one unit of NIIN 0001 is added to authorized stock. To understand the last column in Table 6, recall that the return to a unit of an item is the reduction in the value of the pipeline from stocking that unit, divided by the unit cost; that is,

$$\begin{aligned}&\text{reduction in value of pipeline per dollar} \\ &\qquad = \{\text{WRA inductions per day}\} \times \{\text{unit price of WRA}\} \\ &\qquad\quad \times \{\text{reduction in } E(AWP)\}/\{\text{unit price of item}\}.\end{aligned} \tag{4.3}$$

This example considers a single WRA, so the factor {WRA inductions per day} × {unit price of WRA} is the same for all four items. Thus, the four items are differentiated by {reduction in E(AWP)}/{unit price of item}, which appears in the last column in Table 6. Table 6 indicates which of the items among those going on this WRA should be stocked first: 0002, because it offers the greatest return, i.e., the greatest reduction in value of pipeline per dollar invested in parts.

We have now determined the first entry in the rank-ordered list of increments to authorized stock for this WRA. By going through the above procedure repeatedly (iteratively), we could build the rest of the rank-ordered list of increments to stock for this WRA, obtaining a list like that in Table 7. (The first line in Table 7 is from the sample computation given above; the remaining lines are notional but typical.)

Table 7 illustrates several features of our stockage method. The algorithm stocks items one unit at a time, so it may be advantageous to stock several units of a certain item before stocking any units of other items. (For really cheap items, such as washers, our

method accommodates increments to stock of more than one unit [see MR314].) Also the algorithm often, but not always, selects the cheap parts first, as shown in Table 6.

Table 7

Rank-Ordered List of Increments to Authorized Stock of Parts Going on Hypothetical WRA

NIIN	Unit Price	E(AWP) Before (days)	E(AWP) After (days)	E(AWP) Reduction/$
0002	$ 4	11.17	10.03	0.285
0002	4	10.03	9.00	0.257
0002	4	9.00	8.35	0.162
0001	400	8.35	4.35	0.010
0003	200	4.35	2.35	0.010
0004	50	2.35	1.85	0.010
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.

Doing the Calculation for All the WRAs in a SMIC

In the example just shown, we rank-ordered increments to authorized stock for parts on a single WRA. We can construct similar rank-orderings for each WRA in a SMIC. However, if we want to merge these lists into a single rank-ordered list for the SMIC, we cannot use {reduction in E(AWP)}/{unit price of piece part}, as we did in Table 6. Instead, we must use the full objective function given in Section 3, the reduction in the value of the pipeline from stocking that unit, divided by the unit cost, from Eq. (4.3).

If we were to do this for the hypothetical WRA in Table 6, we would multiply each item in the last column by {WRA inductions per day}×{unit price of WRA}. If we did the analogous computation for each WRA in the SMIC, we would then be able to construct a single rank-ordered list of increments to authorized stock for all the WRAs in the SMIC.

Table 8 shows some results from applying our stockage method to all 11 SMICs that we considered. It displays the first two repair parts selected within each SMIC, along with the unit price of the selected repair parts. The returns in the right-hand column, as defined in Eq. (4.3), are extremely high because this list of increments to stock was built from a basis of zero authorized stock. As one goes down the list of increments to stock within any SMIC, the returns decline quickly to less than "stratospheric" values.

CUT OFF THE LIST TO DETERMINE THE STOCK TO BE AUTHORIZED

We have illustrated how the methods from Section 3 can be used to construct rank-ordered lists of increments to authorized stock, and we have constructed a separate list for each SMIC. Consider the list for SMIC BE. To identify authorized stockage levels for parts going on WRAs of SMIC BE, the list of increments to stock must be cut off at some point since resources are not unlimited; then all of the increments above the cutoff are added to authorized stock.

This cutoff can be made in several ways:

1. By dollar value—e.g., the total amount to be invested in repair parts for BE is $100,000—in which case the cutoff would be made by going down the list until the total of the unit prices of the stocked units was $100,000.

2. By stocking until the average AWP time across WRAs in BE was acceptable.
3. By stocking until the return reached some predetermined value, e.g., a 2:1 return.
4. By reaching a certain depth of stockage within the SMIC, e.g., take the first 1,000 units of increments to stock within each SMIC.

In Section 5, we used a cutoff that was a combination of the second and fourth methods. This is not a suggested policy, but rather a way to exercise the computational tests.

Table 8

First Two Repair Parts Within 11 SMICs

SMIC	IIC	WRA Unit Price	NIIN	Unit Price	E(AWP) days Before	E(AWP) days After	WRA Induction Rate	ROI[a]
BE	E2T0	$ 12710	008828045	$1	23.6	17.9	0.0557	$4007
	E9U9	167050	LLM304417	1	33.2	31.3	0.0082	2554
BP	D761	88290	009445757	3	51.7	51.6	0.9881	4822
	D1D9	48860	LLND02304	2	30.8	30.6	1.1324	3566
CY	KTW6	303050	000035339	1	33.7	33.6	0.9324	14362
	KTW6	303050	004985733	1	33.6	33.6	0.9324	5810
DH	EQP3	106300	000165855	1	34.6	34.5	0.1406	1588
	EQP3	106300	009084990	6	34.5	34.2	0.1406	740
DQ	NWY8	41800	009788570	1	17.8	17.3	0.1763	3428
	LJ64	41800	009788570	1	18.4	17.8	0.1105	2947
EQ	NVN1	45530	009035046	1	13.8	13.6	0.4037	3578
	JP92	46520	009033628	1	16.0	15.8	0.3671	3319
FQ	HA66	57340	003901851	1	31.9	31.7	0.2603	3521
	K522	97930	009138842	1	29.0	29.0	1.0256	1713
MH	ACR4	5500	009007457	1	17.9	4.2	0.0082	620
	NN58	156630	LLND84153	2	34.1	34.0	0.2347	241
PQ	P4Y2	144900	011291032	1	20.2	19.8	0.4785	24747
	P4Y2	144900	011288209	1	19.8	19.5	0.4785	26909[b]
RA	K5T1	83190	009577817	2	14.7	13.9	0.0283	882
	K5T1	83190	008893116	1	13.9	13.6	0.0283	757
TN	Q7V6	103960	011691751	1	34.4	34.3	0.5416	3280
	QKK4	72720	011460174	1	31.7	31.5	0.1900	2904

[a]ROI can be expressed as a function of other columns of the table: ROI = {WRA unit price} × {E(AWP | after) − E(AWP | before)} × {WRA induction rate}/{unit price}. Values are not exact because of rounding.

[b]This second item selected has a higher return than the first item. The ROI is *not* an error: It can happen because the algorithm we use to select increments to stock does not find the optimal solution, but a heuristic solution—albeit a pretty good one.

5. EVALUATING THE METHOD OF BUILDING AUTHORIZED STOCKS

In the preceding section, we discussed how we used Navy data to construct authorized stockage levels of repair parts for NADEPs, for 11 SMICs with large numbers of repair jobs. Table 8 shows some sample results that began with zero authorized stocks and selected the first few units of authorized stock for each SMIC. In this section, we discuss a more realistic stockage exercise and evaluate the resulting stocks to check our method.

The more realistic stockage exercise was as follows:

- We constructed a baseline case by imputing "current" authorized stockage levels at NADEPs, as described below.[26]
- Then we constructed a "treatment case" by adding increments to the baseline stockage levels, using the method described in Section 3.

We evaluated our stockage method by computing the difference in performance between the baseline stockage levels and the treatment stockage levels

- As "predicted" by our stockage method
- As assessed by two simulation tests, a nonparametric test and a parametric test, described below.

In other words, our method predicts certain reductions in AWP time and the repair pipeline as consequences of specific investments in repair parts. The simulation tests assess whether our method does predict the savings that follow from the investments in repair parts. These simulation tests are approximations to actual NADEP operations, but are nevertheless tests that our method must pass to be considered furtl. r.

CONSTRUCTING BASELINE AND TREATMENT CASES

Baseline Case: Imputing Current Stockage Levels

We imputed stocks of repair parts separately for each of the 11 SMICs we used. This subsection describes for a single SMIC how we did the imputation. The objective was to construct authorized stockage levels for repair parts that yield AWP times like the *actual* AWP times in our data set.[27] To build such authorized stockage levels, we began with zero stockage levels for all repair parts, including SRAs, and built a long, rank-ordered list of increments to authorized stock, as discussed in Section 4. We cut off the list so that the resulting authorized stockage levels produced, according to the nonparametric simulation test described below, an average AWP time (across the SMIC) close to the actual average AWP time observed in our NIMMS data. We used the nonparametric test because it is the most realistic alternative available to us.

From the above process we achieved a baseline case. The SMIC average AWP times in the actual data set (and according to the nonparametric simulation test) are in Table 9, along with the number of jobs for each SMIC and the number of repair-part NIINs in the imputed current stock.

Whereas the baseline does not allow us to go as far as claiming that "the Navy will realize X gain if it invests Y dollars," it does allow us to run our simulations with real data and to see whether the tall-pole computation does a reasonable job at estimating the return on the margin from investments in repair parts. Again, the most desirable baseline would be current actual stockage levels at NADEPs, which we do not have and which may not even

exist. The baseline we just constructed has the deficiency that it does not reflect real stockage levels or the discipline under which new stocks are ordered.

Table 9

Aggregate Characteristics of Imputed Baseline Stocks

SMIC	Description	No. of Actual Jobs	No. of Parts, imputed stock	Actual Avg AWP	Avg AWP, imputed stock[a]
BE	E2/C2 electronic a/c	4,526	2,328	35.8	34.6
BP	P3 patrol aircraft	7,915	3,865	23.3	22.6
CY	AWG-9 radar	3,408	1,424	29.7	31.0
DH	H3 helicopter	1,556	491	33.6	34.3
DQ	T56 engine	6,048	3,347	11.3	10.6
EQ	T58 engine	9,055	2,592	16.4	14.9
FQ	T64 engine	3,591	1,971	22.0	16.4
MH	H46 helicopter	2,524	1,421	25.3	25.7
PQ	TF30 engine	6,515	2,914	13.1	12.8
RA	A6E attack aircraft	3,466	2,409	18.4	17.1
TN	F404 jet engine	5,158	2,353	32.8	33.2
ALL		53,762	25,115	21.9	21.0

[a] Computed using nonparametric simulation test.

Treatment Case: An Additional Investment in Stock

Recall that in building the baseline case, we constructed (for each SMIC) a long rank-ordered list of increments to authorized stock. When we set up the baseline case, we cut off that long list at the point where the average AWP time was roughly equal to the observed average AWP time. To build the treatment case, we went down the list another 1,000 increments to authorized stock.

This investment in stock adds the same number of units (1,000) for every SMIC.[28] It does not add the same dollar value to authorized stocks of parts for every SMIC, nor does it cause the same reduction in average AWP time for each SMIC[29]—and therefore would not be a wise policy to implement. Still, as with the baseline case, it enables us to assess whether our method predicts marginal reductions in AWP time.

THE TWO SIMULATION TESTS

We constructed two types of simulation, which we call the parametric test and nonparametric test. Each test involves computations wherein we step through days of the year and

1. For each day, induct WRA jobs at historical rates.
2. For each job, simulate parts requisitions.
3. For each requisition, simulate requisition fills: if the part is authorized and in stock, the requisition is filled in zero days; otherwise, look to the due-in pipeline; otherwise, order the part from the supply system.
4. For both types of simulation, order replenishment stock when the stock on hand falls below authorization levels, using an (S − 1, S) order rule.[30] Table 10 summarizes the types of inputs for the two tests. Following it, we describe the tests in greater detail.

Table 10

Inputs for the Two Types of Tests

Type of Input	Parametric	Nonparametric
1. Job arrival period	1 year	3 years
2. Number of jobs	(No. in NIMMS)/3	(No. in NIMMS)
3. Days of job arrival	Equally spaced	Day of first requisition
4. Repair-parts demands	Binomial simulation parameters	Actual requisitions
5. Requisition-fill times	0 if in stock; negative exponential otherwise	0 if in stock; actual times if asked supply system; negative exponential otherwise

The Parametric Test

The objective of the parametric test was mainly to assess whether the mathematical approximation inherent in the tall-pole calculation introduces serious errors. This test mimics the theory underlying our method.

The parametric test simulated the workings of a NADEP with given authorized stockage levels over the course of a year. For each WRA, the job induction rate is the same as that in our NIMMS data, although the simulated jobs arrived at a constant rate (i.e., at equal deterministic time intervals). Binomial sampling determined which parts had to be replaced. If replacement parts had to be requisitioned from the supply system, OSTs were stochastic draws from negative exponential distributions. For a given SMIC and COG, the mean of the negative exponential distribution for OST was taken from Table 4. The simulation explicitly represented parts on the shelf at the NADEP, consumption of parts for jobs, and the due-in pipeline. Otherwise, the simulation conformed to the assumptions given in Section 3; in particular, it assumed that all parts requirements for a job are determined and ordered on induction.

The Nonparametric Test

Like the parametric test, the nonparametric test simulated the operations of a NADEP and the performance of a given set of authorized stockage levels. However, the nonparametric test was somewhat more ambitious than the parametric test: Its goal was to test whether the tall-pole computation and the other assumptions of our method hold up when confronted with the irregularity of actual job-induction times, actual OSTs, etc.

The nonparametric simulation used the actual times of job inductions as given in NIMMS for the three-year period of our data set. More precisely, for each job, it used the date of the job's first requisition. It also used the actual parts requisitioned for each job. As before, we assumed that all parts requirements for a job were determined and ordered on induction.

Our intention in the nonparametric simulation was to interpose less math between the actual data and the simulation than in the parametric simulation. The distinctive feature of the nonparametric simulation is the way it represents OST and the effect of the authorized stock. Two steps were involved in building this representation:

1. We used NIMMS data to set service times for all requisitions for all WRA jobs in the SMIC, assuming authorized stockage levels of zero for all repair parts.
2. We altered those service times to obtain the times that jobs wait for parts under the authorized stockage levels being assessed.

The purpose of Step 1 is to use the data to say what each OST would have been if the NADEP had kept no stock at all. For nonstocked items (in our NIMMS data, these were all repairables and nonstocked consumables), this step is straightforward: just use the requisition service time in NIMMS, which describes supply system performance for that requisition.[31]

For stocked items, it is not so straightforward to say what would have happened if the NADEP had not been allowed to keep stock. For such items, the requisition service times in NIMMS reflect not just supply system performance but also the beneficial effects, i.e., overall shortening of repair cycle, of whatever stocks the NADEP was allowed to keep. We cannot see, from NIMMS data, how long it would have taken to get such parts had their authorized stockage levels been zero.[32] So we imputed OSTs for these requisitions, using the same model for OST that was used in the parametric simulation: negative exponential with mean determined by SMIC and COG from Table 4.

For Step 2, consider a particular repair part, say, a circuit card. The NIMMS data give us a sequence of requisitions to the supply system for that circuit card, placed on dates on which demands for the card occurred. Our approach is simply to shift the order date back to the first previous requisition date for that part if one item of stock is authorized, to the second previous requisition date if two items are authorized, etc., with the earliest requisitions being shifted back to day one. Appendix B makes this point more explicitly by illustrating a hypothetical sequence of demands.

Discussion of Parametric and Nonparametric Tests

Obviously, these simulation tests do not replicate reality, although the nonparametric test is the more realistic of the two. But the nonparametric test still has several unrealistic features, the most important being that it assumes that all parts required for a job are known on the date the job is inducted and that those parts are all needed on the day of induction. Thus, its measure of AWP time differs from the actual measure of AWP time. Also, the nonparametric test is weakened because, for NIF-stocked items, we had to impute requisition times under zero stockage, instead of using times directly drawn from NIMMS. Nonetheless, the two simulation tests, particularly the nonparametric test, do stress our method because they deviate from its assumptions in the direction of reality.

SIMULATION RESULTS

We ran the simulations described above. Here we present the results, summarizing the effects for the nonparametric and parametric cases at both aggregate (SMIC) and disaggregate (WRA) levels.

Summary of Nonparametric Test—Aggregate Level

Table 11 shows the value of the repair pipeline with zero stock, with current stock, and with the stock prescribed in our baseline and treatment allocations. Current stocks are unknown, but we do not need to know what they are to compute the value of the repair pipeline.[33]

Table 11

Value of the Repair Pipeline at Different Stockage Levels, Nonparametric Test

	Value of the Repair Pipeline				Cost of Repair Parts	
	Stockage Level ($ millions)				Stockage Level ($ thousands)	
SMIC	Zero	Current	Baseline	Treatment	Baseline	Treatment
BE	7.5	4.8	3.2	2.2	180	347
BP	11.9	8.3	6.3	4.8	372	620
CY	19.1	14.5	15.2	10.2	19	255
DH	1.8	1.6	1.4	0.8	2	74
DQ	2.5	0.9	0.7	0.5	302	496
EQ	4.8	3.0	1.6	1.0	189	363
FQ	9.5	3.6	2.3	1.4	130	316
MH	4.2	3.6	1.8	1.1	291	815
PQ	10.7	5.8	4.0	2.8	279	572
RA	4.9	3.3	3.2	2.6	194	378
TN	9.8	5.4	4.2	2.9	290	637
TOTAL	86.7	54.8	43.9	30.3	2,248	4,873

We estimate that, with zero stock, the AWP part of the repair pipeline for the above 11 SMICs would have had $86.7 million worth of inventory.[34] Current stockage reduces that value to $54.8 million (a reduction of 36.8 percent). Our baseline stockage reduces that value to $43.9 million (a reduction of 49.4 percent). Recall that our baseline allocation mimics actual average wait times within SMIC, but it weights acquisitions according to their reduction in pipeline value. To the extent that current stocks do not quantify an attempt to reduce pipeline value, one would expect to see pipeline values for the baseline-stockage position that are smaller than what occurs in the current system. We estimate that our baseline-stockage position reduces the pipeline value by 20 percent compared with the current-stockage position ($43.9 million versus $54.8 million).

Allocating money to stocking parts will be a wise decision only if, for every dollar invested in parts, there is an anticipated reduction of one dollar or more in pipeline value. The return on investment (ROI) ratio captures this quantity:

$$\mathrm{ROI} = \frac{\begin{array}{c}\text{value of repair pipeline}\\ \text{(initial stockage level)}\end{array} - \begin{array}{c}\text{value of repair pipeline}\\ \text{(with additional stock)}\end{array}}{\text{cost of additional stock}}. \tag{5.1}$$

When one starts with zero stock, computed savings are large. For the baseline system, we constructed $2.2 million in stock. Thus, our simulated return on investment (SROI) ratio for increments to zero stock is 19.5 ([86.7–43.9]/2.2). But one also expects decreasing returns on investment, which is the purpose of the treatment case. Comparing the baseline and treatment cases suggests a further decline of about $13.6 million for a further investment of $2.6 million, or an SROI ratio of 5.2 ([43.9–30.3)]/2.6).

Table 12 displays various aspects of the pipeline reductions in our treatment and baseline cases. It shows differences in pipeline values versus differences in costs of repair parts, sorted by SROI. It demonstrates that the tall-pole computations predicted the pipeline reductions quite well at the SMIC level. The correlation of the tall-pole return on investment (TpROI) and SROI is 0.98. The correlation is still high (0.94) if the highest-return SMIC—

CY—is omitted. In summary, whenever the tall-pole calculation suggests that investing in a SMIC is a high-return case (TpROI ≥ 2, meaning that for every $1 invested in parts there is a $2 reduction in pipeline value), the simulation agrees (SROI ≥ 3).

Table 12

Value and Cost Summaries, Sorted by Simulated Return on Investment, Nonparametric Test

SMIC	Treatment Cost ($ thousands)	Pipeline Reduction ($ millions) Simulation	Pipeline Reduction ($ millions) Tall-Pole	Return on Investment Ratio Simulation	Return on Investment Ratio Tall-Pole
CY	236	5.0	4.2	21.0	17.9
DH	72	0.6	0.7	9.3	9.7
BP	248	1.5	1.3	6.2	5.3
BE	167	0.9	0.6	5.6	3.8
FQ	185	1.0	0.5	5.2	2.6
PQ	292	1.2	0.7	4.2	2.5
TN	347	1.3	1.2	3.8	3.4
RA	184	0.6	0.6	3.4	3.1
EQ	174	0.5	0.3	3.0	2.0
MH	524	0.7	0.6	1.4	1.2
DQ	195	0.1	0.1	0.8	0.5
ALL	2,624	13.4	10.8	5.2	4.2

It is important to emphasize that these SROIs apply to the simulation, not necessarily to the real world. The strongest reason not to extrapolate beyond the simulation is that we do not know actual stocks. Although we made a good-faith effort to divert attention from the incredibly large SROIs that would occur in making the baseline stockage zero, we have no solid basis to say that the ROI would be 5.2. We conclude that any further attempts to validate these calculations in the real world must start with knowledge of actual stocks, as well as knowledge of the system by which new stocks are ordered.

Summary of Nonparametric Test—Disaggregate Level

The above observations pertain to the SMIC level. Nine of the 11 SMICs yielded high returns. A natural question to ask is, Might we expect to see the same results at the WRA level? That is, when we stock for a WRA, do we see returns for that WRA, or do we just see high-return WRAs on average, with relatively little ability to differentiate among WRAs?

At the WRA level, however, we have to deal with randomness. From the description of the nonparametric test, recall that random draws enter into the calculations when, for example, we need to requisition the supply system for a part. At the SMIC level, results are somewhat insensitive to randomness because there are many jobs, and we are taking averages, so the law of large numbers eliminates much of the random variation.

In fact, variances in WRA SROIs go down as 1/n, where n is the number of inductions of a WRA in the simulation. Figure 7 displays box-and-whisker plots[35] of results for WRAs that enter the simulation many times (no. of jobs ≥ 45, which is the 66th percentile). For these plots, the WRAs were divided into 10 groups based on their TpROIs. The medians of each group are shown below the horizontal axis. Summary statistics for the SROIs are represented by the box-and-whisker plots. They display the median (middle; white line), 25th and 75 quartiles (outer box), 1.5 times interquartile range (whiskers), and outlying values (horizontal lines). As the medians for each of the ten groups increase, the summary statistics show correspondin; increases. Thus, the relationship between TpROI and the SROI is direct. Also, there is little enough *dispersion* (the height of the box) that what looks like a high-return WRA (according to the tall-pole) usually is one.

Figure 8 shows a similar picture for WRAs that enter the simulation infrequently (no. of jobs ≤ 11, which is the 33rd percentile). We still see a direct relationship by noting that, as the medians for the TpROIs increase, so do the medians for the SROIs. But, because each WRA has few inductions, the simulated returns show enough dispersion about the mean that many WRAs indicated by the tall-pole calculation to be good investments actually turned out in the simulation to be poor investments. A substantial portion of their box is below 1, meaning that more money was spent stocking parts than was removed from the repair pipeline.

RAND#146-7-0194

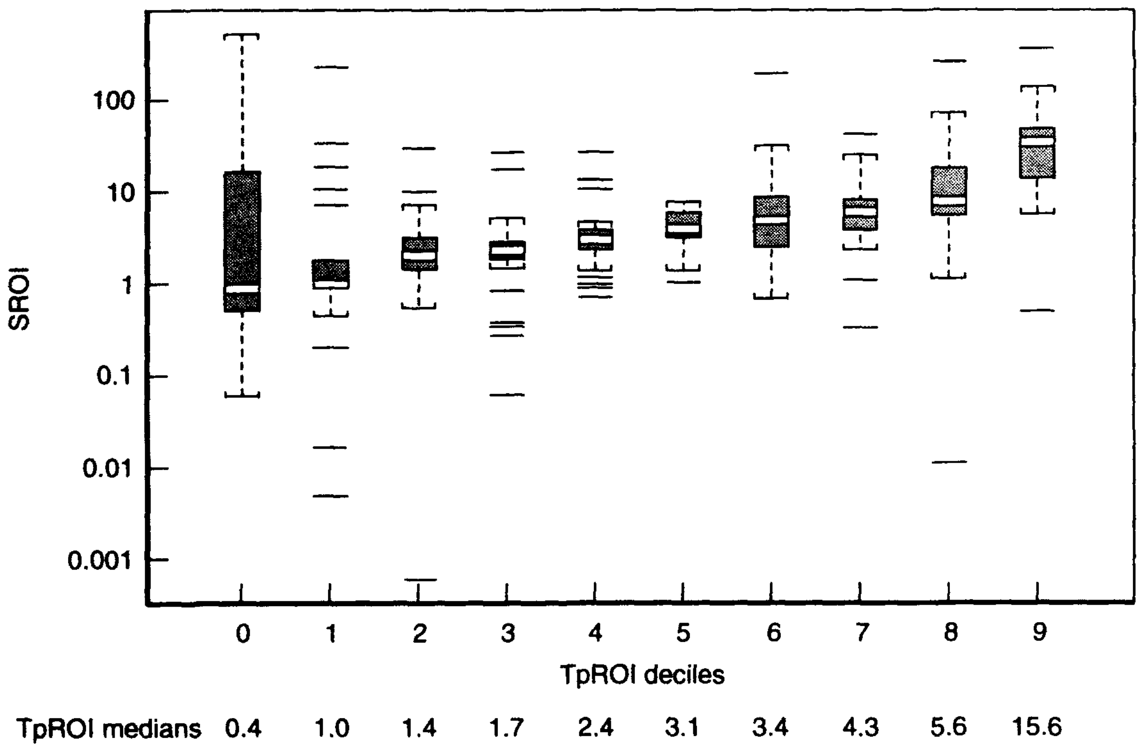

Figure 7—Box-and-Whisker Plots of WRA Tall-Pole ROIs Versus Simulated ROIs—Frequent Jobs (≥45), Nonparametric Simulation

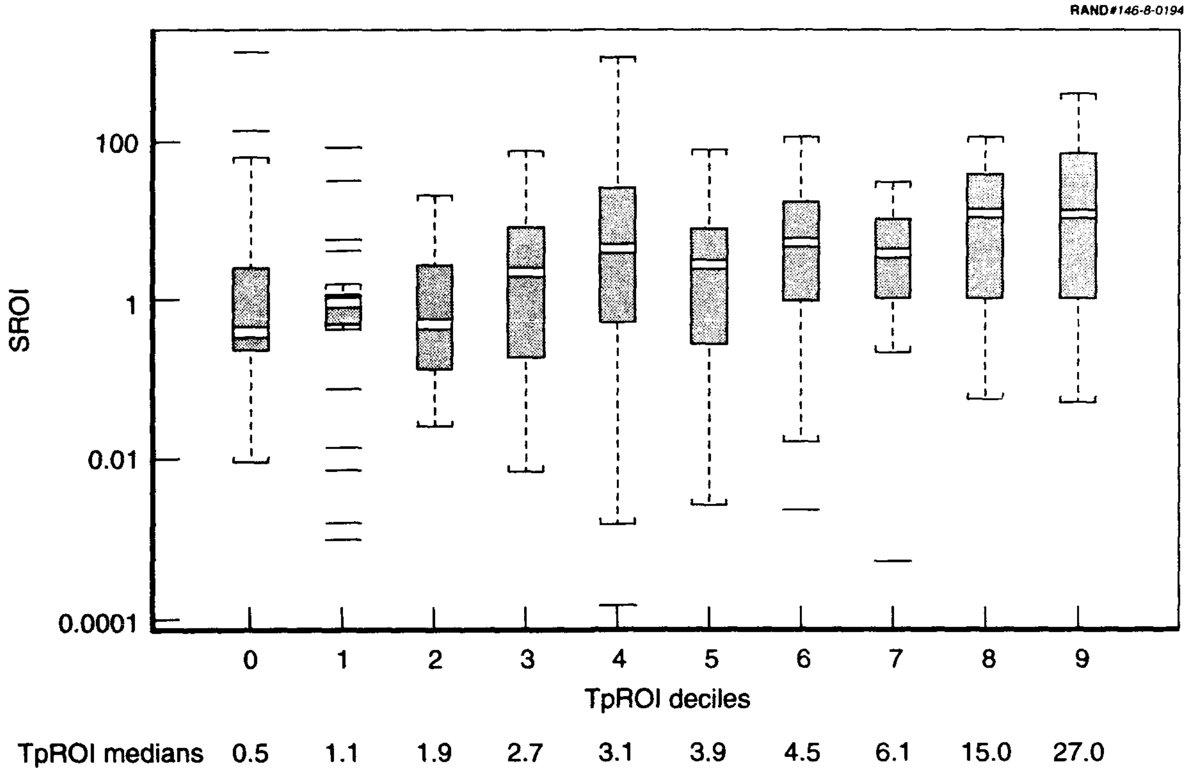

Figure 8—Box-and-Whisker Plots of WRA Tall-Pole ROIs Versus Simulated ROIs—Infrequent Jobs ($\leq$ 11), Nonparametric Simulation

We attempted to summarize this variation numerically for a set of interesting cases: TpROIs ≥ 2 (i.e., cases for which the tall-pole indicated that the investment would likely pay off). We ran a weighted regression of the logarithm of the simulated ROIs on the logarithm of the tall-pole ROIs. The weights were the numbers of inductions over the course of the year. The regression fit is

$$\log_{10}(\mathrm{SROI}) = 0.13728 + 0.95043 \times \log_{10}(\mathrm{TpROI}), \quad (5.2)$$

with a standard deviation of 0.47509. The coefficient of 0.95 signifies that the SROI is close to being proportional to the TpROI, but the large standard deviation indicates that we cannot predict the returns with much accuracy when the number of inductions is small.

Table 13 summarizes fitted values and variation about fitted values for selected TpROIs and for number of inductions ranging from one per year to one per day. The fitted value is derived from Eq. (5.2): fitted value = $10^{\log_{10}(\mathrm{SROI})}$. The table also displays the exponentiated $\pm$ 2 standard deviation intervals around the fitted values, which ought to capture about 95 percent of the SROIs. The table shows how the width of those intervals varies according to the rates of induction. As the frequency of induction increases, the width of the interval decreases, signifying that there is less variation around the fitted value. Infrequently inducted items (one per year) have yields that are unpredictable. The -2 standard deviation limit in fitted SROI is often less than 1.0, meaning that for every dollar invested in stocking parts there is less than a dollar reduction in pipeline value. Not every investment is going to pay off: It takes the law of large numbers to justify investments in

stocking parts. On the other hand, 12 inductions per year seem to be enough; most WRAs inducted that often look like sure high-return WRAs.

Table 13

WRA-Level Summary: Nonparametric Simulation ROIs Versus Tall-Pole ROIs

Tp-ROI	IndRate (per year)	Fitted Value	-2 SD Limit	+2 SD Limit	Tp-ROI	IndRate (per year)	Fitted Value	-2 SD Limit	+2 SD Limit
2	1	2.7	0.3	22.6	10	1	12.2	1.4	104.4
	4		0.9	7.7		4		4.2	35.8
	12		1.4	4.9		12		6.6	22.7
	52		2.0	3.6		52		9.1	16.5
	365		2.4	3.0		365		10.9	13.7
5	1	6.3	0.7	54.0	20	1	23.6	2.8	201.8
	4		2.2	18.5		4		8.1	69.1
	12		3.4	11.8		12		12.7	43.9
	52		4.7	8.5		52		17.6	31.8
	365		5.7	7.1		365		21.1	26.5

NOTE: IndRate is induction rate.

The bottom line for the nonparametric simulation is that, after allowing for stochastic variation in the simulations, the TpROI approximates the SROI pretty well; i.e., we saw no strong indication that the tall-pole's approximations and assumptions introduce distortions into the actual evaluation.

Summary of Parametric Test—Aggregate Level

Table 14 shows the value of the repair pipeline with zero stock and with the stock prescribed in our baseline and treatment allocations. This simulation estimates that, with zero stock, the repair pipeline for the above 11 SMICs would have had $95.0 million worth of inventory (versus $86.7 million for the nonparametric simulation). The baseline-stockage position has a pipeline value of $35.9 million (versus $43.9 million), and the treatment case reduces that value to $21.7 million (versus $30.3 million), for a reduction of $14.2 million (versus $13.6 million).

We should not be surprised to see differences in pipeline values. Numerous differences between the two simulations could account for them (see Table 10). What is interesting, however, is the similarity in SROIs. The overall SROI for the parametric simulation is 5.4 (versus 5.2 for the nonparametric simulation). Moreover, the relationship between SROIs and TpROIs within SMICs is about the same. Table 15 displays pipeline reductions between baseline and treatment cases versus differences in costs of repair parts, sorted by SROI. As before, we see that the tall-pole computations predict the pipeline reductions quite well at the SMIC level. The correlation of the TpROI and SROI is 0.9956. The correlation is still high (0.9870) if the highest-return case, CY, is omitted. Whenever the tall-pole calculation suggests that investing in a SMIC yields a high return (TpROI ≥ 2), the simulation agrees (SROI ≥ 2).

Summary of Parametric Test—Disaggregate Level

The above observations pertain to the SMIC level. Again, we ask, Might we expect to see the same results at the WRA level, where randomness is more a factor because of the smaller sample sizes?

Table 14

Value of the Repair Pipeline at Different Stockage Levels, Parametric Test

	Value of the Repair Pipeline				Cost of Repair Parts	
	Stockage Level ($ millions)				Stockage Level ($ thousands)	
SMIC	Zero	Current	Baseline	Treatment	Baseline	Treatment
BE	7.9	4.8	2.5	1.7	180	347
BP	14.4	8.3	4.5	2.4	372	620
CY	21.6	14.5	16.9	11.1	19	255
DH	1.8	1.6	1.3	0.3	2	74
DQ	2.4	0.9	0.4	0.3	302	496
EQ	5.2	3.0	0.9	0.4	189	363
FQ	9.1	3.6	0.7	0.3	130	316
MH	4.9	3.6	0.8	0.2	291	815
PQ	11.4	5.8	2.1	1.2	279	572
RA	6.0	3.3	3.5	2.9	194	378
TN	10.3	5.4	2.3	0.9	290	637
TOTAL	95.0	54.8	35.9	21.7	2,248	4,873

Table 15

Value and Cost Summaries, Sorted by Simulated Return on Investment, Parametric Test

		Pipeline Reduction ($ millions)		Return on Investment Ratio	
SMIC	Treatment Cost ($ thousands)	Simulation	Tall-Pole	Simulation	Tall-Pole
CY	236	5.8	4.2	24.6	17.9
DH	72	1.0	0.7	13.9	9.7
BP	248	2.1	1.3	8.5	5.3
BE	167	0.8	0.6	4.8	3.8
TN	347	1.4	1.2	4.0	3.4
RA	184	0.6	0.6	3.3	3.1
PQ	292	0.9	0.7	3.1	2.5
EQ	174	0.5	0.3	2.9	2.0
FQ	185	0.4	0.5	2.2	2.6
MH	524	0.6	0.6	1.1	1.2
DQ	195	0.1	0.1	0.5	0.5
ALL	2,624	14.2	10.8	5.4	4.2

Figure 9 displays the results for WRAs that enter the simulation many times (no. of jobs ≥ 15, which is the 66th percentile). Again, we provide box-and-whisker plots of SROI in base-10 logarithms. There is a direct relationship between TpROI and the SROI, and little enough dispersion so that what looks like a high-return WRA (according to the tall-pole) usually is one. Figure 10 shows a similar picture for WRAs that enter the simulation infrequently (no. of jobs ≤ 4, which is the 33rd percentile). We still see a direct relationship, but because each WRA has few inductions, the simulated return shows enough dispersion about the mean that many WRAs the tall-pole indicates are good investments turned out to be poor ones.

We again attempted to summarize this variation for cases that the tall-pole thought were high-return investments (TpROI ≥ 2). We ran a weighted regression of log actual gains on the log tall-pole predictions. The weights were the numbers of inductions over the course of the year. The regression fit is

$$\log_{10}(\text{SROI}) = -0.39789 + 1.2613 \times \log_{10}(\text{TpROI}), \qquad (5.3)$$

with a standard deviation of 0.91540. The coefficient of 1.26 signifies that the SROI is close to being proportional to the TpROI (however, not as close as the nonparametric simulation), but the very large standard deviation again indicates that we cannot predict the returns with much accuracy.

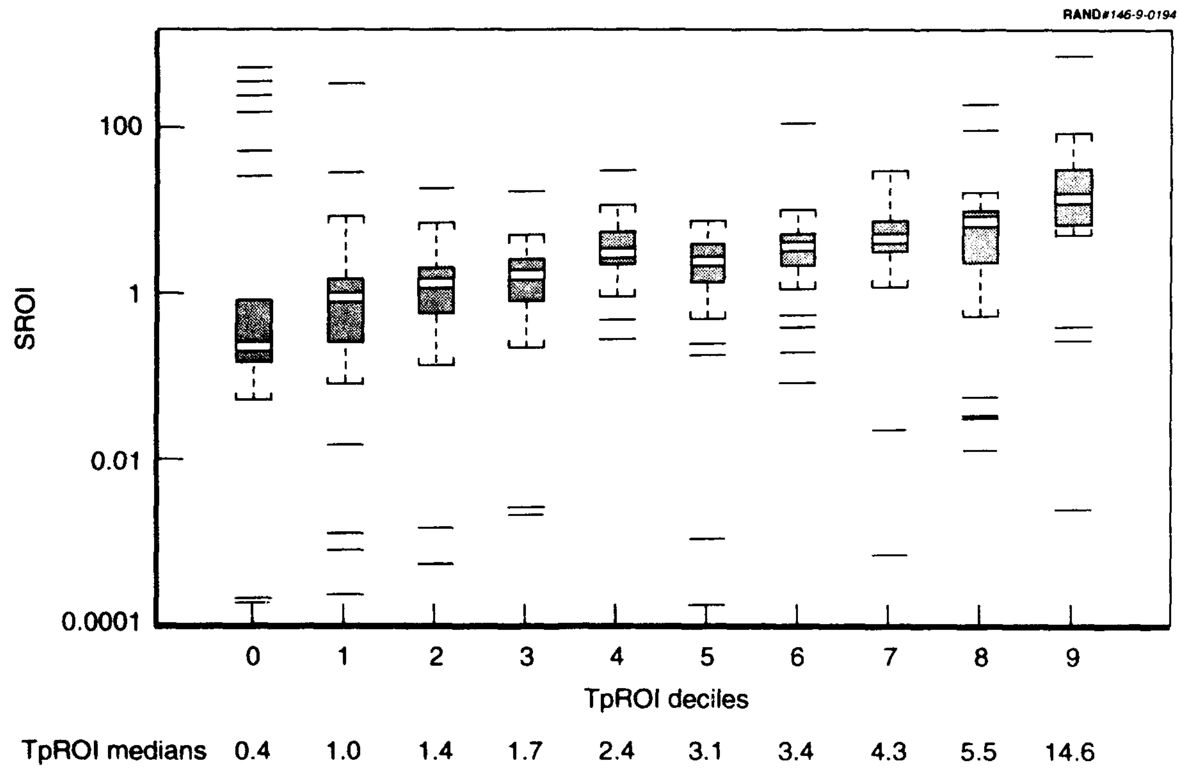

Figure 9—Box-and-Whisker Plots of WRA Tall-Pole ROIs Versus Simulated ROIs—Frequent Jobs (≥ 15), Parametric Simulation

Table 16 summarizes fitted values and variation about fitted values for selected TpROIs and for numbers of inductions going from one per year to one per day. Infrequently inducted items (one per year) had yields that varied considerably around their TpROI values. The −2 standard deviation limit in fitted SROI is often less than 1.0. As before, not every investment is going to pay off: The law of large numbers is needed to justify these investments. On the other hand, for frequently inducted WRAs (once a month or more), the variation in realized SROI is small. If the tall-pole indicated that an investment would pay off, it usually would.

The bottom line for the parametric simulation is the same as for the nonparametric simulation. After allowing for stochastic variation in the simulations, we conclude that the TpROI approximates the SROI pretty well; i.e., we saw no strong indication that the tall-pole's approximations and assumptions introduce distortions into the actual evaluation.

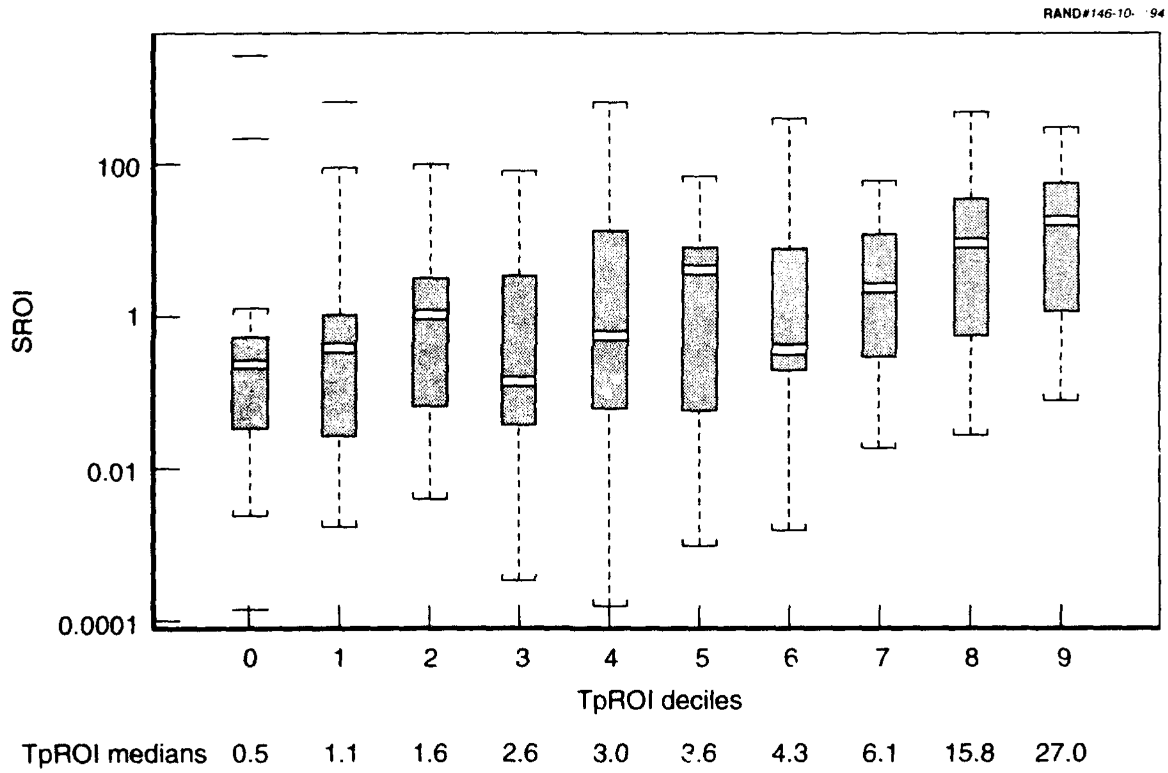

Figure 10—Box-and-Whisker Plots of WRA Tall-Pole ROIs Versus Simulated ROIs—Infrequent Jobs (≤ 4), Parametric Simulation

Table 16

WRA-Level Summary: Parametric Simulation ROIs Versus Tall-Pole ROIs

Tp-ROI	IndRate (per year)	Fitted Value	−2 SD Limit	+2 SD Limit	Tp-ROI	IndRate (per year)	Fitted Value	−2 SD Limit	+2 SD Limit
2	1	1.0	0.0	59.7	10	1	7.3	0.1	454.6
	4		0.1	7.6		4		0.9	57.6
	12		0.3	3.2		12		2.2	24.1
	52		0.5	1.7		52		4.1	12.9
	365		0.8	1.2		365		5.9	9.1
5	1	3.0	0.0	189.6	20	1	17.5	0.3	1,089.6
	4		0.4	24.0		4		2.2	138.1
	12		0.9	10.0		12		5.3	57.7
	52		1.7	5.4		52		9.9	31.0
	365		2.5	3.8		365		14.1	21.7

NOTE: IndRate is induction rate.

6. ADDITIONAL EVALUATION OF THE REPAIR PROCESS

Section 5 describes the simulation programs used to evaluate authorized stockage levels (ASLs). But setting ASLs is not the only lever available for reducing the repair pipeline. For example, the civilian manufacturing industry is shifting away from inventories to just-in-time delivery systems,[36] in which parts arrive from suppliers as they are needed in the manufacturing process. Thus, the emphasis is on predictably low order-and-ship times.

Other RAND studies have shown much interest in examining the effects of responsive distribution. In earlier work for the Navy, we estimated the effect of faster distribution on mission capability.[37] An actual case study of the effects of reduced OST on mission capability in the Air Force's Coronet Deuce Test, as well as a simulation study, analyzed the effects of shortened OST on both mission capability and the level of spares necessary to sustain the two-level maintenance concept being implemented. The simulation study of F-16 two-level maintenance for avionics showed that the cost of spares decreased by $40 million when OST was cut in half and aircraft capability was held constant.[38]

This motivated us to study the effects of reducing OSTs in the NADEP repair environment. We used the parametric simulation method described in Section 5 to examine the potential payoff from reducing OSTs. We asked, How will reducing OSTs affect the value of the repair pipeline? To examine this effect, we reduced the OSTs from the parametric simulation runs described in Section 5 in the manner described below.

SETTING NEW PARAMETERS FOR OSTs

Our starting point for this analysis is Table 4, which lists the average OST for each weapon system (SMIC) and each supply point (COG). We note four very long average times: 120 days or more. Rather than simply reducing all OSTs across the board, we targeted the long times and left the shorter ones alone, by setting four thresholds—2, 7, 14, and 30 days—for the arrival of parts, and varying a shrinkage factor for OSTs beyond those thresholds. The formula is

$$\begin{aligned} \text{new OST} &= \text{OST} && \text{if OST} \leq \text{T} \\ &= \text{T} + (\text{OST} - \text{T}) \times \text{M} && \text{if OST} > \text{T}, \end{aligned} \tag{6.1}$$

where T is the threshold (2, 7, 14, or 30 days) and M is a multiplier (0, .25, .50, .75, 1). Thus, if Table 4 shows a mean OST of 8 days, the new OST with a threshold of 2 days and a 0.5 multiplier would yield $2+(8-2) \times .50 = 5$ days. Note that when the multiplier is zero, the new OST is the smaller of threshold value and the old OST. When the multiplier is 1, the new OST equals the old OST. The other multipliers yield a range of values between the threshold and the old OST.

Using the new mean OSTs and the parametric simulation model described in Section 5, we simulated the demand for parts needed in the NADEP repair pipeline.

THE EFFECT OF REDUCED OSTs ON THE VALUE OF THE REPAIR PIPELINE

The purpose of the simulation was to illustrate how much savings in the repair pipeline might be effected by changing OSTs. The aggregate results of the simulations are shown in Figure 11. Each panel corresponds to a different threshold—from left to right the threshold values are 2, 7, 14, 30 days; within panels, the multiplier values are 0, .25, .50, .75, 1. On the plots themselves, a circle represents the value of the repair pipeline when the

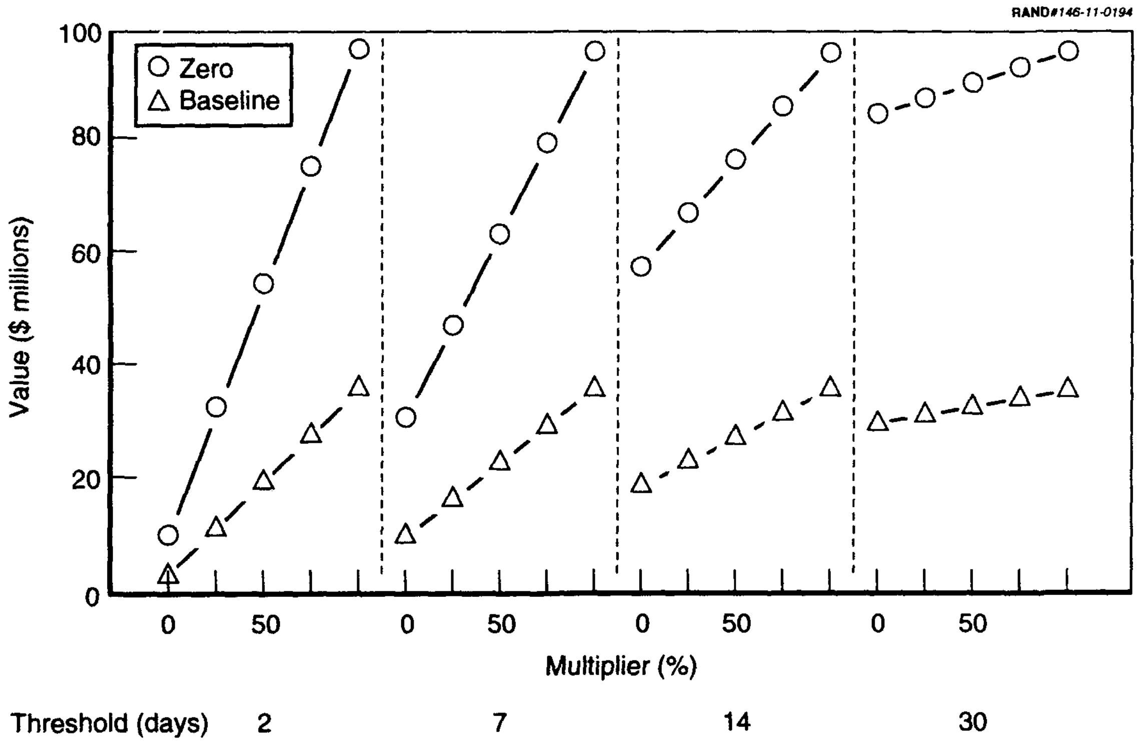

Figure 11—The Total Effect of Reducing OST on the Value of the Repair Pipeline Under Two Different Stockage Postures

initial stockage position is zero; the triangle denotes the initial baseline-stockage position described in Section 5. The vertical axis shows the total value of the repair pipeline, which ranged from $3.2 to $95.0 million, depending on stockage level and OST/multiplier combination.

The main lesson to be learned from these plots is that faster transportation and distribution can do as much as authorized stock to reduce the value of the repair pipeline. Figure 11 shows that if there is no stock in the system and the current OSTs are used in the simulation, the value of the repair pipeline is $95.0 million. When the threshold is 7 days and the multiplier is 0, the pipeline value decreases to $30.2 million, which is less than the baseline-stockage pipeline value ($36.0 million).

We have a measure of the amount of authorized stock—the actual cost of the stock. The corresponding measure for the transportation and distribution computations would be the cost or savings of achieving faster transportation and distribution, but we have no information on that. Instead, we calculate a measure of the average transportation and distribution time under the 20 different scenarios we used in Figure 11. We tabulate their values in Table 17, and plot repair pipeline values against average transportation and distribution times in Figure 12.

The value of the repair pipeline is essentially a linear function of the average OST for repair parts. With zero authorized stock, the correlation between the two is 0.9986, and the slope of that line is $4.0 million per day; i.e., a day's reduction in the transportation and distribution pipeline is worth $4.0 million, under a zero-stock scenario. Under the baseline-stockage scenario, the correlation is still high—0.9956—but a day's reduction is worth only $1.6 million.

Table 17

Aggregate Values of the Repair Pipeline Versus Mean OSTs

Threshold (days)	Multiplier	Repair Pipeline Value ($ millions) Zero Stockage	Baseline Stockage	Mean OST (days)
2	1	95.0	36.0	22.9
2	0.75	74.9	27.1	17.6
2	0.50	54.0	18.8	12.4
2	0.25	32.1	10.7	7.2
2	0	9.7	3.2	2.0
7	1	95.0	36.0	22.9
7	0.75	79.6	29.0	18.9
7	0.50	63.0	22.2	14.9
7	0.25	46.5	15.8	10.9
7	0	30.3	9.9	7.0
14	1	95.0	36.0	22.9
14	0.75	85.4	31.2	20.6
14	0.50	75.5	27.1	18.3
14	0.25	65.5	22.9	16.1
14	0	56.8	19.4	13.8
30	1	95.0	36.0	22.9
30	0.75	92.6	34.6	22.4
30	0.50	89.6	33.0	22.0
30	0.25	86.8	31.7	21.6
30	0	84.2	30.4	21.2

Looking at the formula (from Section 3) for the value of the repair pipeline,

$$\text{total AWP contribution to value of repair pipeline} = \Sigma\,\{\text{unit price of the WRA}\} \times E(\text{WRA inductions per day}) \times E(\text{days of AWP time for the WRA}), \tag{6.2}$$

the preceding result should not be too surprising, because days of AWP time is a multiplier in the pipeline value computation. Faster transportation and distribution matter much more when authorized stockage levels are zero because every job must then reach into the supply system for every part.

As a side calculation, we ran the tall-pole algorithm with the reduced OSTs to see how the list of authorized stock might change. As OSTs were reduced, we expected to see greater range; i.e., different types of items stocked, as opposed to multiple occurrences of the same item. Our reasoning is that multiple demands on the same item are unlikely to occur closely in time, and with low OSTs and (S – 1, S) ordering, the stock is replenished quickly whenever a demand occurs. Table 18 summarizes the range of items at three stages of authorized stocks: when there are 250, 500, and 1,000 items per SMIC. The table shows the number of different stock items per 100 authorized items as the threshold decreases to 2 days with the multiplier M=0. As transportation times get faster, the range increases from about 80 per 100 to over 90 per 100.

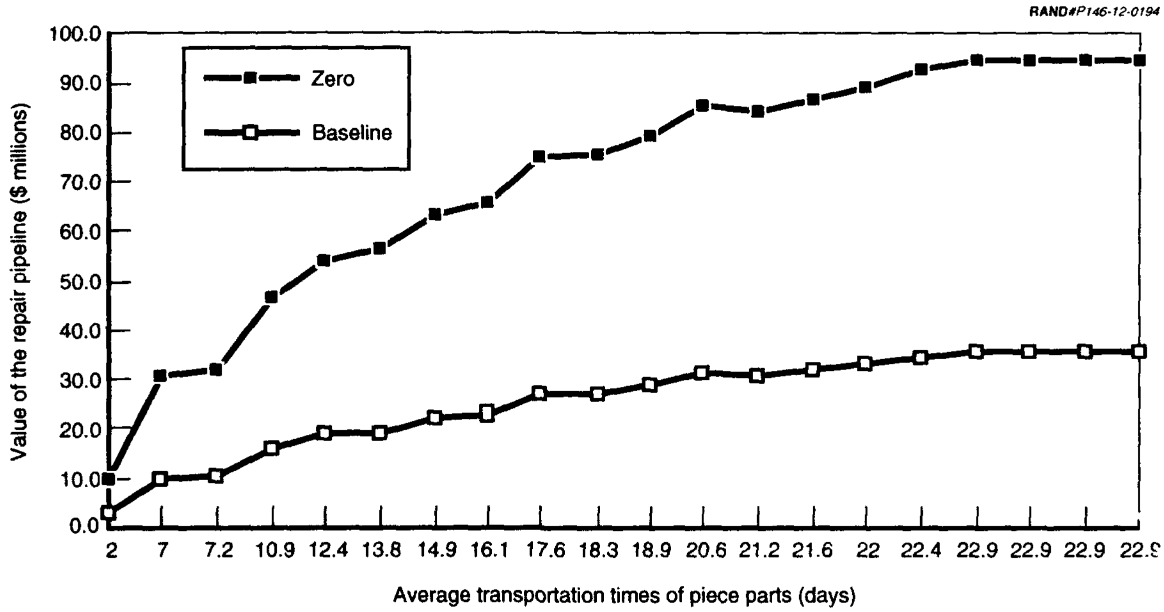

Figure 12—Value of the Repair Pipeline Versus Average Transportation Times of Repair Parts

In summary, AWP time reductions can be achieved in three ways: setting authorized stockage levels, reducing OSTs, or doing both at the same time. Setting ASLs requires an investment in stock, whereas reducing OSTs requires an improvement in process. In either case, the model we developed of the repair process can be used to judge the effects of the supply action.

Table 18

Range of Authorized Stock Lists (number of different items per 100)

	Threshold (days)				
Items per SMIC	2	7	14	30	None
250	95.2	90.4	86.8	85.3	85.5
500	93.5	87.4	83.0	81.6	81.5
1,000	92.4	84.9	79.1	76.3	76.0

7. CONCLUSIONS

Current stockage policies emphasize descriptors of parts (unit price, failure rates), but they rarely include information about the end-item that needs them (the repair of which is affected by their availability). This lack contributes to long repair turnaround times. Inexpensive parts can delay expensive end-items for long periods because the supply system does not know of their importance.

We have developed a method here to help manage the flow of repairable parts through the NADEP. Our method follows from a view of the repair pipeline that models end-item flows and that captures delays due to parts shortages. Our method produces lists of parts to be ordered proactively. Such ordering attempts to reduce total end-item wait times and hence results in more end-items in ready for issue status. We developed mathematical approximations to determine such lists, and, in two types of simulation tests, we showed that the approximations do a reasonable job.

Our model incorporates five important measures: cost of the end-item, demand for the end-item, replacement factor for the repair part, cost of the repair part, and OST required to obtain the repair part. The method combines these measures into a quantity that we call the *value of a part*: dollar cost (in end-items) that a given part takes out of the repair pipeline. The method also quantifies the amount of end-item time saved per dollar invested in repair parts. These measures can be useful in summarizing the flow of parts moving through the NADEP and in assessing the impact of alternative policies designed to increase NADEP throughput (and hence decrease the amount of money tied up in the repair pipeline).

RESULTS

Our research suggests that, through effective stockage of repair parts, the Navy may be able to achieve large savings. Those savings come from shortening turnaround time at the NADEP, which then allows more end-items to be in circulation. Furthermore, our evaluations suggest that our calculations can identify weapon systems for which it would make sense to stock more parts and those for which it would not. The calculations can be used to balance investment strategies between spending money on parts and spending it on other segments of the repair pipeline.

The next step is to collect more accurate data about the repair process, extrapolate the effect of new stockage policies, and evaluate their effectiveness both in terms of cost reduction and, perhaps, increase in mission capability. This evaluation could be done through raw calculations or experimentation at a NADEP.

NIMMS provided enough information to construct candidate lists of parts. But to evaluate the real savings from candidate lists, current inventory positions, bills of material, order-and-ship times, WRA demands, etc., must be known. To project actual savings would require additional data acquisition and validation. Although we were unable to do this, we did run simulations using the NIMMS data and employing assumptions consistent with those data.

The simulations showed us several things. First, the approximations worked. Second, stocking repair parts and using $(S - 1, S)$ ordering could do much to reduce the value of the repair pipeline. Third, with the data sets we constructed, it was possible to predict some high-return and low-return SMICs. Fourth, because of random variation, some high returns and low returns can be expected at the WRA level, but the law of large numbers prevails;

that is, investments pay off mainly in the aggregate, and by particular WRA for frequently inducted WRAs.

We primarily assessed the effects of incrementing the list of authorized stock, but we also assessed the effects of reduction in OST. Table 4 displays average OSTs as shown in the NIMMS data, by COG and by SMIC, some of which are longer than four months. We ran the parametric simulations for both the zero- and baseline-stockage positions and 16 different OST parameterizations. For the baseline case, with a 50 percent cut in OST and a threshold of 7 days, the value of the repair pipeline is $22.2 million, down from $35.9 million, which compares favorably with the value of $21.7 million obtained by acquiring an additional 1,000 items for each SMIC, at a cost of $4.9 million (from Table 14). To complete the evaluation, one would have to know the costs of reducing OSTs by the specified amounts, and to compare them with the costs of an investment in additional stock (e.g., the $4.9 million) that would achieve the same repair pipeline value.

POSSIBLE USES OF THE VALUE MEASURE

Finally, we should point out that there are several other ways to use the value of parts idea:

1. Should the Navy decide to stock parts at NADEPs the way it stocks parts on deploying carriers in AVCALs (Aviation Consolidated Allowance Lists), our method could serve as a starting point for building those stocks. The Navy has called this idea a *NADEPCAL*.
2. Value of parts could be used to evaluate current reorder rules for existing stocks and to identify how those rules should change as weapon systems age, the force is drawn down, NADEP repair procedures change, etc.
3. The calculations can be used to balance investment strategies between spending money on parts and spending it on other segments of the repair pipeline: parts reliability, faster requisition processing, reduced repair times, faster distribution, etc.
4. The calculations can be used to identify parts that are problematic or that may become problematic.
5. A similar computation of value measure can be used to attribute value to supply actions (e.g., speedup of delivery of due-in items) in terms of the effect such actions have on the availability of aircraft at the end of a specific time horizon.

This report has concentrated on using a value measure to make investments in repair parts to minimize the expected value of the repair pipeline for a given level of investment; in addition, as suggested above in the fifth point, a similar computation of value can be used to help managers make operating decisions that affect the short-run output of the NADEP.[39] With knowledge of the near-term repair schedule, the stock on hand, and the parts due in from supply, managers at NADEPs can take actions to alter the parts due in to service the repair schedule more expeditiously. Using a value measure in this way provides the NADEP with a tool that indicates when it is worthwhile to selectively pay more for better supply service (for example, using Federal Express), targeting expenditures to specific items that will yield the largest improvement in aircraft availability. We have not had the opportunity to test the short-run value measure as we tested the long-run measure.

Because the Navy was interested only in reducing depot throughput, the model is single echelon.[40] The multi-echelon dimension, opportunistic use of parts, and cannibalization (controlled substitution) are refinements to be added to the model. One of

the biggest distinctions between our model and others currently in use is the modeling of multiple replacements—multiple repair parts are frequently involved in one repair action, so that if parts 1, 2, and 3 are needed and only parts 1 and 3 are in stock, the repair is still delayed. We also note that data limitations precluded extending the computations to arbitrary levels of indenture, even though the mathematics for such computations are described in MR314.

Experiences in the private sector show that changes in both policies and processes are necessary to achieve order-of-magnitude improvements in remanufacturing operations. In addition to setting authorized stock levels for the NADEPs, the Navy must give attention to improving maintainability and contracting, updating replacement factors and demand rates, and decreasing OST and process times. We conclude that assessing these various options through their effect on the value of the repair pipeline is a good way to proceed.

NOTES

[1] *Second-destination transportation* refers to all transportation of parts and equipment after the service has received the item initially. *First-destination transportation* is from the original manufacturer to the service.

[2] The NADEPs at Norfolk, VA; Pensacola, FL; and Alameda, CA, are being closed, leaving the Navy with NADEPs at North Island, CA; Cherry Point, NC; and Jacksonville, FL.

[3] During the period covered by our data, DLA managed simple, cheap parts such as bolts and washers, and the Navy's Aviation Supply Office (ASO) handled more sophisticated parts, such as circuit cards and actuators, some of which were consumable and others of which were repairable. All parts purchased by the Navy are identified as *consumable*, meaning that they are discarded when they break, or *repairable*, meaning that someone tries to fix them when they break. Under the new arrangement, DLA will be responsible for consumables and ASO, for repairables. The Navy's concern is that DLA is more removed from the repair processes than is ASO.

[4] NADEPs prefer to maintain a steady workload, so for parts that are frequently repaired, e.g., radar antennas, level schedules are negotiated every six months. These schedules help the NADEPs plan for personnel, parts, equipment, etc. For example, in the period of our data, NADEP Norfolk had a level schedule of 12 AWG-9 radar antennas per quarter.

[5] See, for example, Lionel A. Galway, *Management Adaptations in Jet Engine Repair at a Naval Aviation Depot in Support of Operation Desert Shield/Storm*, Santa Monica, Calif.: RAND, N-3436-A/USN, 1992, or Lionel A. Galway, *Materiel Support Problems at a Naval Aviation Depot: A Case Study of the TF-30 Engine*, Santa Monica, Calif.: RAND, N-3473-A/USN, 1992.

[6] We assume that a replenishment requisition for part i is placed after every s^{th} demand for part i. This is often called (S – s, S) ordering, where S is the authorized stockage level. If S = 0—the part has no authorized stock—then we set s to 1 and place an order after each demand. See, for example, S. M. Ross, *Stochastic Processes*, New York: Wiley, 1983, p. 69.

[7] The *repair pipeline* is composed of all the events that occur from the time the WRA is removed from the aircraft on the flight line until it is placed on the shelf at the local supply center in RFI status.

[8] We will use WRA as the end-item example throughout the document. However, the method has a much wider application in any remanufacturing setting, such as when diesel engines or hydraulic lifts are overhauled.

[9] The replacement factor is defined for each repair part. Because not all parts on a WRA are needed for each repair, the *replacement factor* for a part is the proportion of repair jobs that require the part.

[10] The *bill of materials* is a list of all the parts on the WRA and their indentured relationship. A WRA is made up of piece parts and SRAs. An SRA is also made up of piece parts and sometimes of other SRAs. The distinction between WRAs/SRAs and piece parts is that the former can be repaired and the latter cannot.

[11] Galway, *Management Adaptations in Jet Engine Repair*, 1992.

Galway (p. 7) also noted, "Many Navy jet engines either have no BOM or have one that is limited in scope. This is *not* to say that artisans do not know what parts are typically needed for repair jobs. However, since repair parts are never requisitioned in advance of engine induction and preliminary inspection, little attention has been given in recent years to centrally recording BOM and repair factor information. This topic is controversial and is under reconsideration within the depot system."

[12] Hodges, James S., "NIF Stockage Policies: An Annotated Briefing," Santa Monica, Calif.: RAND, unpublished research, and Hodges, James S., and J. Payne, "An Adaptive Approach to Material Support: Work in Progress," Santa Monica, Calif.: RAND, unpublished research.

[13] Although 40 days may seem long, for many combinations of inventory control points (ICPs) and weapon systems (actually, SMICs, special materiel identification codes), the average customer wait time at depots is 40 days or greater (see Table 4). Reducing average OST may be a higher-payoff strategy than stocking one unit of an item.

[14] See, for example, George B. Dantzig, *Linear Programming and Extensions*, Santa Monica, Calif.: RAND, R-366-PR, 1963, pp. 517–520.

[15] The greedy algorithm derives its name from the fact that each additional part is chosen to give maximal value. See, for example, Eugene L. Lawler, *Combinatorial Optimization: Networks and Matriods*, New York: Holt, Rinehart, and Winston, 1976, pp. 275–280.

[16] See, for example, S. M. Ross, *Introduction to Stochastic Dynamic Programming*, New York: Academic Press, 1983, Chapter V.

[17] A particular stockage policy is characterized by its *depth* (how many of a particular part are stocked) and its *range* (the number of different items stocked).

[18] See, for example, S. M. Ross, *Stochastic Processes*, New York: Wiley, 1983, p. 69.

[19] For these assumptions, and assuming that in each WRA or SRA job, each part is demanded independently with constant probability, each part defines an $M/M/n/\infty$/FIFO (First In, First Out) queue. See, for example, D. Gross and C. M. Harris, *Fundamentals of Queueing Theory*, Wiley: New York, 1974.

[20] The tall-pole algorithm is an unsophisticated approximation. It was forumlated in the spirit of Occam' razor, a scientific and philosophic rule that stresses simplicity over complexity when choosing competing theories.

[21]A *common item* is a part that is used on more than one weapon system, such as a circuit card that goes on both the AWG-9 radar and the APG-70 radar.

[22]Military depots perform two distinct types of repair. *Component repair* is the repair of WRAs and SRAs that have broken in the field. *SDLM* is regularly scheduled aircraft maintenance and repair at a level not available in the field.

[23]In terms of NIMMS record codes, we included requisitions initiated by code 11 or code 12 records and excluded direct issues (code 32) or back-ordered requisitions (code 14).

[24]The *cognizance code* is a symbol used exclusively within the Navy to identify the ICP that exercises supply management over the item. The meanings for all the COGs are listed in Appendix A. The most frequently occurring codes in the data are from only a few ICPs listed below.

COG	ICP	Description
1R	ASO	Consumable aeronautical material
7R	ASO	Repairable aeronautical material
9F	Warner-Robins ALC	Consumable aeronautical material
9I	Ogden ALC	Consumable aeronautical material
9J	Okla. City ALC	Consumable aeronautical material
9K	Sacramento ALC	Consumable aeronautical material
9Z	Defense Industrial Supply Center (DISC)	Consumable material

NOTE: ALC is Air Force Air Logistics Center.

[25]We do not assert that get times are nonstochastic. Rather, we assume it in order to drastically simplify our computations. We depend on the simulation studies of Section 5 to establish that the resulting approximations are useful.

[26]We would like to have used actual NADEP authorized stocks. However, we do not have data on what they are. Thus, we had to invent "current" stockage levels, which, obviously affects the interpretation of our evaluation. We discuss the effect later in this section.

[27]Recall that our definition of *AWP time* assumes that as soon as a WRA is inducted, the parts needed to repair it are known and that they are ordered immediately. In reality, our actual NIMMS data gave us customer wait times.

[28]Adding 1,000 items to each SMIC is not presented as a good policy for the Navy to adopt; rather, it was a way to check the tall-pole approximation against the simulations. It produces a big spread in rates of return among SMICs, which reflects why it is not a good policy.

[29]These would have been additional ways of assessing the effectiveness of the method.

[30]Note that with an $(S - 1, S)$ reorder rule, every time a part is requisitioned, a replacement part is automatically ordered.

[31]We did so for all requisitions initiated by a code 11 or code 12 record. Almost all of those requisitions were Issue Priority Group I or II.

[32]We did not want to use service time for replenishment requisitions, because such requisitions are mostly Issue Priority Group III and thus quite a bit slower than requisitions for nonstocked items.

[33]Recall the formula

$$\begin{aligned}\text{value of repair pipeline} = \sum &(\text{WRA cost})\\ &\times(\text{WRA induction rate})\\ &\times(\text{average AWP times for the WRA}),\end{aligned}$$

which does not directly involve stockage levels.

[34]This computation replaces OSTs for stocked items with OSTs used in the parametric simulations, i.e., it draws from a negative exponential distribution with mean determined by SMIC and COG from Table 4.

[35]These box-and-whisker plots were produced through the statistical software package *S-Plus*. Richard A. Becker, John M. Chambers, and Allan R. Wilks, *The New S Language: A Programming Environment for Data Analysis and Graphics,* Pacific Grove, Calif.: Wadsworth & Brooks/Cole, 1988.

[36]Cummins Engine was able to reduce its inventory from $173 million to $22 million and Detroit Diesel reduced safety stock from 30 days' supply to 5 days' by implementing responsive ordering, transportation, processing, and holding operations.

[37]Brauner, Marygail K., Daniel A. Relles, and Lionel A. Galway, *Improving Naval Aviation Depot Responsiveness*, Santa Monica, Calif.: RAND, R-4133-A/USN, 1992.

[38]Abell, John B., and H. L. Shulman, *Evaluations of Alternative Maintenance Structures*, Santa Monica, Calif.: RAND, R-4205-AF, 1992.

[39]The mathematics that underlie both the long-run and short-run uses of the value measure are provided in the companion document, MR314.

[40]Our model is single echelon because it considers only work performed at the depot. A multi-echelon model would consider work performed at the depot and intermediate level or at the depot, intermediate level, and the flight line.

APPENDIX

A. LIST OF INVENTORY CONTROL POINTS

COG	ICP	Description
0Z	Unknown	—
1H	SPCC	Navy Stock Fund materiel
1R	ASO	Consumable aeronautical materiel
5R	ASO	Consumable aircraft launch and recovery equipment
6K	ASO	Repairable aeronautical materiel assigned to containers
6V	NADEP	Technical directive change kits
7R	ASO	Repairable aeronautical materiel
9A	Navy Fleet Materiel Support Office (NFMSO)	Parts peculiar to combat vehicles
9C	Defense Construction Supply Center	—
9D	Defense Personnel Support Center	—
9F	Warner-Robins ALC	Consumable aeronautical materiel
9G	Defense General Supply Center	General defense materiel
9I	Ogden ALC	Consumable aeronautical materiel
9J	Okla. City ALC	Consumable aeronautical materiel
9K	Sacramento ALC	Consumable aeronautical materiel
9L	NFMSO	Defense medical materiel
9N	NFMSO	Electronic materiel
9Q	General Services Admin.	General support materiel
9S	Army Missile Materiel Readiness Command	Consumable materiel
9V	Sacramento ALC	Consumable materiel
9W	Army Troop Support and Aviation Materiel Readiness Command	Consumable aviation materiel
9X	Defense Fuel Supply Center	Petroleum materiel
9Y	Army Communications and Electronics Materiel Readiness Command	Consumable materiel
9Z	Defense Industrial Supply Center (DISC)	Consumable materiel

NOTE: ALC is the Air Force Air Logistics Center; SPCC is the Navy Ships Parts Control Center.

B. TREATMENT OF DEMANDS FOR NONPARAMETRIC TEST

A hypothetical sequence of demands is represented in Table B.1. The dates of the demands are in the leftmost column; because zero stockage was assumed, the units to fill the demands are requisitioned from the supply system on the same day the demands occur. The second column gives the number of each requisition, and the number in the third column gives the date each requisition was placed. The next two columns are the OST (either taken directly from the NIMMS data or constructed, as described in Section 5) and the date the unit arrived. These data imply the length of time each demand waited under the assumption of zero stockage, which is in the sixth column from the left.

Table B.1

Illustration of Nonparametric Simulation for Hypothetical Circuit Card

	Stockage Level Zero					Stockage Level One				
Date of Part Demand	Req. No.	Date Unit Ordered	OST	Date Unit Arrived	Wait	Date Unit Ordered	OST	Date Unit Arrived	Wait	Change in Wait
10	1	10	2	12	2	1	2	3	0	–2
20	2	20	5	25	5	10	5	15	0	–5
30	3	30	8	38	8	20	8	28	0	–8
40	4	40	5	45	5	30	5	35	0	–5
50	5	50	30	80	30	40	30	70	20	–10

Now suppose that the authorized stockage level for the circuit card is one unit. Suppose also, following our model in Section 3, that (S – 1, S) ordering is used for the circuit card, so that each time a unit is demanded for a repair job, a replacement is ordered from the supply system. The effect of having an authorized stockage level of one unit is that each demand for the circuit card is filled by a unit that was ordered from the supply system after the previous demand.

The key to the nonparametric simulation is to alter the data for "stockage level zero" to mimic that effect of authorized stockage. Here is how we did that. Under a stockage level of zero, the demand on date 50 was filled by requisition no. 5; under a stockage level of one, the demand on date 50 is still filled by requisition no. 5, but requisition no. 5 is placed earlier, on day 40, the date of the previous demand for the circuit card. We assume that requisition no. 5 still takes 30 days for the supply system to service, but because it was placed on day 40 instead of day 50, the unit arrives on day 70 instead of day 80, 10 days earlier than under the assumption of zero authorized stockage. Similarly, the unit demanded on day 40 is serviced by requisition no. 4; but with an authorized stockage level of one, requisition no. 4 was placed on day 30 instead of day 40, so it is also serviced 10 days earlier than it would be if the stockage level were zero. The seventh through tenth columns show what happens with a stockage level of one. The reduction in the amount of time that the demands wait to be filled is in the rightmost column.

By analogy with the argument just given, if the authorized stockage level of the circuit card is two units, the unit used to fill a given demand was ordered from the supply system after the second previous demand. If the authorized stockage level is two units, then, the

demand at time 50 is serviced by requisition no. 5, which is placed on day 30 and arrives on day 60. And so on. By the same kind of argument, any authorized stockage level can be accommodated.

Once the service times ("wait" in Table B.1) have been altered to reflect the authorized stockage levels, the AWP time for a job is then determined as the maximum of the wait times for the parts on that job.